

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 IMPLICIT NEURAL COMPRESSION OF POINT CLOUDS VIA LEARNABLE ACTIVATION FUNCTION

Anonymous authors

Paper under double-blind review

ABSTRACT

Efficiently compressing and transmitting large-scale high-fidelity 3D point clouds is a critical bottleneck for practical applications. We introduce a novel framework that reformulates point cloud compression as model compression. Our framework models high-fidelity point cloud geometry and attribute with compact implicit neural representations (INR) separately and then compresses the model parameters directly via quantization and entropy coding, decoupling representation from compression. To ensure this neural representation is both faithful and efficient, we employ Kolmogorov-Arnold Network (KAN) as the INR backbone. Thanks to its superior approximation properties and parameter efficiency, KAN can easily capture fine-grained details missed by traditional MLP. Extensive evaluations on datasets such as KITTI, ScanNet, and 8iVFB demonstrate that our method significantly outperforms the MPEG standard and prior implicit neural representation approaches. Notably, it achieves competitive rate-distortion performance against state-of-the-art deep learning codecs. Our findings establish implicit neural compression as a powerful and practical pathway for developing the next generation of high-efficiency point cloud codecs.

1 INTRODUCTION

Point clouds have emerged as a foundational data modality for 3D perception, powering critical applications in autonomous driving (Li et al., 2020; Cui et al., 2021), augmented and virtual reality (AR/VR) (Lim et al., 2022; Wang et al., 2023), and embodied intelligence (Qi et al., 2024). The proliferation of advanced LiDAR sensing technologies (Raj et al., 2020) has made the acquisition of large-scale, high-resolution point clouds more feasible than ever.

However, raw point clouds' massive scale, spatial sparsity, and lack of explicit topological structure create substantial memory and bandwidth overhead, which severely impedes their practical deployment (Graziosi et al., 2020). Consequently, the development of efficient Point Cloud Compression (PCC) solutions becomes a critical and pressing necessity.

Early efforts to standardize PCC were developed by the Moving Picture Experts Group (MPEG), resulting in the release of two foundational frameworks: geometry-based PCC (G-PCC) (Schwarz et al., 2018) and video-based PCC (V-PCC) (Graziosi et al., 2020). Deep learning has subsequently driven a paradigm shift in the field. Many of these methods (Huang & Liu, 2019; Quach et al., 2020; Que et al., 2021; Zhang et al., 2024) are built upon generic encoder-decoder architectures, which encode a point cloud into a compact latent representation for subsequent reconstruction. While these data-driven codecs often surpass the rate-distortion (RD) performance of traditional standards, their dependence on large-scale pre-training datasets limits generalization to out-of-distribution data.

To address the generalization limitations of data-driven methods, a new paradigm based on implicit neural representations (INR) has emerged (Xue et al., 2024; Ruan et al., 2024a;b). INR-based methods do not directly learn how to reconstruct point clouds. Instead, they train a lightweight implicit neural representation to model the distribution of a single point cloud in 3D space. This strategy inherently avoids the generalization issues of data-driven codecs. However, how to select the optimal INR backbone for compression remains an open research problem, as it must balance the conflicting demands of capturing fine-grained details and achieving a low compression rate.

054 In this work, we introduce **Point cloud Implicit neural COmpression (PICO)**, a framework that re-
 055 formulates PCC from a signal processing problem to a neural network compression problem. PICO
 056 begins by modeling the point cloud’s geometry and attributes using two compact INRs separately.
 057 Subsequently, PICO directly compresses the parameters of these learned INRs through advanced
 058 quantization and entropy coding techniques. This paradigm provides two key advantages. First, it
 059 decouples geometry and attribute modeling, avoiding the detrimental feature entanglement. Sec-
 060 ond, it separates point cloud representation from compression, enabling fine-grained control over
 061 the compression rate and reconstruction quality.

062 PICO incorporates a multi-scale rate control mechanism that allows precise and dynamic bitrate
 063 allocation, providing a notable advantage over existing methods. For coarse-grained control, we
 064 select an optimal model architecture using a pre-computed Pareto frontier that profiles the trade-off
 065 between model size and bitrate. To achieve finer-grained adjustments, we then apply a tunable ℓ_1
 066 regularization during training to promote parameter sparsity. This sparsity facilitates compression
 067 of the trained model, allowing the final bitrate to be precisely determined by adjusting the quanti-
 068 zation step size during entropy coding. Through jointly optimization of model size, sparsity, and
 069 quantization, PICO achieves precise bitstream control while preserving high compression quality.

070 PICO adopts the Kolmogorov-Arnold Network (KAN) (Liu et al., 2025) as its INR backbone in-
 071 stead of the typical multilayer perceptron (MLP). Inspired by KAN, we design a backbone called
 072 **Learnable Activation Function Network (LeAFNet)**. Compared to MLP, its learnable activation
 073 functions can better capture the high-frequency details in point clouds, while achieving comparable
 074 INR performance with fewer parameters, which is a critical factor for compression tasks. To fur-
 075 ther enhance performance, we adapt LeAFNet for PCC by adding positional encoding to improve
 076 spatial understanding and replacing B-spline functions with radial basis functions to increase model
 077 throughput. These modifications make LeAFNet a backbone specifically designed for PCC.

078 PICO improves practical deployability by optimizing sampling space and strategy, as well as in-
 079 troducing dynamic thresholding, which together substantially reduce computational overhead and
 080 memory footprint. In addition, we explore how to extend the PICO from static to dynamic point
 081 clouds, broadening its applicability to a wider range of point cloud types.

082 We evaluated PICO against MPEG standards and other PCC methods on the 8iVFB, KITTI (Geiger
 083 et al., 2013), and ScanNet (Dai et al., 2017a) datasets, and conducted ablation studies to validate the
 084 effectiveness of our design. On the 8iVFB dataset, PICO showed strong performance. It reduced
 085 BD-BR by 53.54% and improved BD-PSNR by 4.92 dB for geometry compression, and in the more
 086 challenging joint compression task, it achieved a 42.71% BD-BR reduction and a 2.70×10^{-3}
 087 improvement in BD-PCQM. These results provide strong evidence of the efficiency of PICO.

088 Our main contributions can be summarized as follows:

- 089 ❶ We propose PICO, an implicit neural PCC framework with precise rate control mechanism, which
 090 is further optimized for real-world deployment and applicable to a wide range of point cloud types.
- 092 ❷ We propose LeAFNet, an INR backbone with learnable activation functions, which is lightweight
 093 and highly effective at fitting implicit functions, making it well-suited for PCC.
- 094 ❸ We reformulate PCC as neural network compression and conduct extensive experiments to assess
 095 its potential as a foundation or component of next-generation PCC.

097 2 RELATED WORK

100 2.1 POINT CLOUD COMPRESSION

101 The MPEG 3D Graphics Coding Group has established the PCC standard, introducing two meth-
 102 ods: G-PCC and V-PCC (Graziosi et al., 2020). G-PCC directly encodes geometry and attributes,
 103 using an octree with entropy coding for voxelized geometry and Trisoup meshes for low-bitrate sur-
 104 face approximation, while attributes are handled by transforms such as region-adaptive hierarchical
 105 transform (RAHT) (De Queiroz & Chou, 2016). V-PCC projects point clouds into 2D patches and
 106 attribute maps via planar parameterization, then packs them into video frames for High Efficiency
 107 Video Coding (HEVC) (Sullivan et al., 2012) compression, using geometry-color separation and
 motion-compensated prediction to maintain continuity.

Recent advances in deep learning have significantly improved PCC with autoencoder-based frameworks, which encode point clouds into compact latent representations and reconstruct them via learned decoders, using entropy models to optimize the rate–distortion trade-off. To address point cloud irregularity, two main architectures have emerged: voxel-based methods (Wang et al., 2021b;a; 2022), which apply hybrid 3D convolutions on volumetric representations, and point-wise methods (Huang & Liu, 2019; Sheng et al., 2021), inspired by PointNet (Qi et al., 2017), which operate directly on raw point coordinates without voxelization artifacts. Although these approaches often outperform traditional codecs in rate–distortion performance, they continue to struggle with generalization to unseen domains and scalability to large-scale scenes (Quach et al., 2022).

To overcome the limitations of encoder–decoder frameworks, recent research has explored a new paradigm for PCC based on INRs (Ruan et al., 2024a; Xue et al., 2024; Ruan et al., 2024b). These methods represent an entire point cloud as a continuous, coordinate-conditioned neural function, which mitigates the generalization limitations of prior learning-based approaches. However, these works do not investigate how to select optimal INR backbone model for target bitrates, nor do they address efficiency considerations for practical deployment.

2.2 IMPLICIT NEURAL REPRESENTATION

INR (Ramasinghe & Lucey, 2022; Saragadam et al., 2023; Sitzmann et al., 2020) parameterize continuous multidimensional signals using coordinate-based neural networks. Given an input coordinate $\mathbf{x} \in \mathbb{R}^d$, a neural network f_θ , whose parameters θ are optimized to minimize the reconstruction error with respect to the ground truth signal $s(\mathbf{x})$, outputs the corresponding signal value $f_\theta(\mathbf{x}) \approx s(\mathbf{x})$. Formally, an INR can be expressed as:

$$f_\theta : \mathbf{x} \in \mathbb{R}^d \mapsto s(\mathbf{x}) \in \mathbb{R}^c, \quad \theta^* = \arg \min_{\theta} \mathcal{L}(f_\theta(\mathbf{x}), s(\mathbf{x})), \quad (1)$$

where \mathcal{L} denotes a suitable reconstruction loss, d is the input coordinate dimension, and c is the signal dimension. INRs have been widely used for data compression in other domains. For instance, COIN and other methods (Dupont et al., 2021; Strümpler et al., 2022; Dupont et al., 2022) map pixel coordinates to pixel colors and use meta-learning to improve fitting efficiency. These works offer valuable insights for developing INR-based PCC.

2.3 KOLMOGOROV-ARNOLD NETWORK

KAN (Liu et al., 2025) is a network architecture designed to improve function approximation and interpretability. Unlike MLP with fixed activation functions, KAN is inspired by the Kolmogorov–Arnold representation theorem, which states that any multivariate continuous function can be expressed as a superposition of univariate continuous functions. Leveraging this insight, KAN decomposes complex high-dimensional functions into combinations of simpler one-dimensional functions, using learnable one-dimensional functions as activation units. This design enables KAN to achieve exceptional representational efficiency, making it particularly effective for modeling high-frequency details and complex signal, and providing a powerful tool for INR-based applications.

3 PICO

In this section, we introduce PICO. We first briefly describe how the two-stage compression is implemented, and then present the improvements PICO makes for the two-stage process. Next, we describe how PICO achieves rate control and the design details of LeAFNet. The pseudocode of the compression and decompression algorithm is provided in detail in Alg. 1 and Alg. 2.

3.1 TWO STAGE COMPRESSION

The 3D point cloud $\mathcal{P} = \{\mathcal{X}, \mathcal{A}\}$ in N -bit voxelized space \mathcal{S} typically comprises two components, geometry \mathcal{X} representing a set of 3D coordinates, where the coordinate \mathbf{x} satisfies:

$$\mathbf{x} = (x, y, z) \in \mathcal{S} = \left\{ \left(\frac{k_x}{2^N}, \frac{k_y}{2^N}, \frac{k_z}{2^N} \right) \mid k_x, k_y, k_z \in \mathbb{Z}, 0 \leq k_x, k_y, k_z < 2^N \right\}, \quad (2)$$

162 and attributes \mathcal{A} representing color, material, or reflectance. **PICO** compresses the point cloud \mathcal{P} by
 163 applying geometry compression to \mathcal{X} and attribute compression to \mathcal{A} .
 164

165 **Stage 1: Geometry Compression.** We train the first INR f_g to learn a continuous occupancy field.
 166 f_g takes a coordinate from \mathcal{S} as input and outputs the probability p that this coordinate is occupied
 167 by the geometry \mathcal{X} . Subsequently, we binarize the continuous field using a threshold τ , marking
 168 each coordinate as either occupied or unoccupied to obtain reconstructed geometry $\hat{\mathcal{X}}$:

$$f_g : \mathbf{x} \in \mathcal{S} \mapsto p \in [0, 1], \quad (3)$$

$$\hat{\mathcal{X}} = \{\mathbf{x} \mid f_g(\mathbf{x}) > \tau, \mathbf{x} \in \mathcal{S}\}. \quad (4)$$

172 **Stage 2: Attribute Compression.** After obtaining $\hat{\mathcal{X}}$, we train a second INR f_a to learn the corre-
 173 sponding attributes. f_a takes a coordinate from $\hat{\mathcal{X}}$ as input and outputs the normalized attribute \mathbf{c}
 174 (consider color as the default attribute). We use $\tilde{\mathcal{A}}$ as the training ground truth, which is obtained by
 175 mapping attributes from \mathcal{A} to $\hat{\mathcal{X}}$ using a nearest neighbor principle. By traversing all coordinates in
 176 $\hat{\mathcal{X}}$ through f_a , we obtain the reconstructed attributes $\hat{\mathcal{A}}$:

$$f_a : \mathbf{x} \in \hat{\mathcal{X}} \mapsto \mathbf{c} \in [0, 1]^3, \quad (5)$$

$$\tilde{\mathcal{A}}(\hat{\mathbf{x}}_i) = \mathcal{A} \left(\arg \min_{\mathbf{x}_j \in \mathcal{X}} \|\hat{\mathbf{x}}_i - \mathbf{x}_j\| \right), \quad (6)$$

$$\hat{\mathcal{A}} = \{f_a(\mathbf{x}) \mid \mathbf{x} \in \hat{\mathcal{X}}\}. \quad (7)$$

184 We consider the INR set $\{f_g, f_a\}$ as a proxy for the compressed point cloud. It is only necessary
 185 to store and transmit $\{f_g, f_a\}$. During point cloud decompression, we can obtain the reconstructed
 186 point cloud $\hat{\mathcal{P}} = \{\hat{\mathcal{X}}, \hat{\mathcal{A}}\}$ by traversing the spatial coordinates through Eq. 4 and Eq. 7.

189 3.2 SAMPLING SPACE AND STRATEGY

191 **Sampling Space.** Due to the inherent sparsity of the original point cloud, the vast majority of
 192 voxels in the voxelized space \mathcal{S} are empty, which makes training and inference on the whole \mathcal{S}
 193 computationally prohibitive and difficult to optimize. To address this issue, we divide the original
 194 space \mathcal{S} into $2^M \times 2^M \times 2^M$ coarse-grained cubes, and **PICO** processes only the set of non-empty
 195 cubes, denoted as \mathcal{W} , during both training and inference. The optimized sampling space \mathcal{V} is thus
 196 defined as the union of all voxels within these non-empty cubes:

$$\mathcal{W} = \{\mathbf{w} \mid \mathbf{w} = \lfloor \mathbf{x} \cdot 2^M \rfloor / 2^M, \mathbf{x} \in \mathcal{X}\}, \quad (8)$$

$$\mathcal{V} = \{\mathbf{x} \mid \lfloor \mathbf{x} \cdot 2^M \rfloor / 2^M \in \mathcal{W}, \mathbf{x} \in \mathcal{S}\}. \quad (9)$$

200 **Sampling Strategy** Although redefining the sampling space from \mathcal{S} to \mathcal{V} eliminates a large num-
 201 ber of empty voxels, non-empty voxels still constitute only a tiny fraction δ within \mathcal{V} . This severe
 202 class imbalance presents a huge challenge for training. To mitigate this problem, we use weighted
 203 sampling to control the proportion of positive labels $\alpha = 0.5$ in each training batch \mathbf{x} . Sampling \mathbf{x}
 204 from both the non-empty voxels \mathcal{X} and empty voxels $\mathcal{V} \setminus \mathcal{X}$ can be expressed as:

$$\mathbf{x} = \alpha \cdot U(\mathcal{X}) \oplus (1 - \alpha) \cdot U(\mathcal{V} \setminus \mathcal{X}), \quad (10)$$

207 where $U(\cdot)$ denotes uniform sampling, \oplus denotes the concatenation. However, computing the empty
 208 voxels $\mathcal{V} \setminus \mathcal{X}$ incurs a time complexity of $O(|\mathcal{V}| \cdot |\mathcal{X}|)$, while storing them requires $O(|\mathcal{V}|)$ memory.

209 To reduce the time overhead, we avoid the costly operation of explicitly generating $\mathcal{V} \setminus \mathcal{X}$. Instead,
 210 We perform approximate sampling separately from the non-empty voxels \mathcal{X} and the redefined sam-
 211 pling space \mathcal{V} . Eq. 10 can then be rewritten as:

$$\mathbf{x} = \hat{\alpha} \cdot U(\mathcal{X}) \oplus (1 - \hat{\alpha}) \cdot U(\mathcal{V}), \quad (11)$$

213 where the calibrated sampling rate $\hat{\alpha} = (\alpha - \delta) / (1 - \delta)$ is used to maintain the target class ratio.
 214 This strategy lowers the time complexity to $O(1)$, resulting in a more efficient sampling process.

216 **Algorithm 1** `PICO_Compression`

217 **Input:** point cloud $\mathcal{P} = \{\mathcal{X}, \mathcal{A}\}$, model dictionary \mathcal{M} , bitrate b , voxel space \mathcal{S}

218 **Parameter:** learning rate γ , cube resolution M , voxel resolution N

219 **Output:** bitstream $\hat{\theta}_g, \hat{\theta}_a$

220 1: $\mathcal{W} \leftarrow \{\mathbf{w} \mid \mathbf{w} = \lfloor \mathbf{x} \cdot 2^M \rfloor / 2^M, \mathbf{x} \in \mathcal{X}\}$

221 2: $\mathcal{V} \leftarrow \{\mathbf{x} \mid \lfloor \mathbf{x} \cdot 2^M \rfloor / 2^M \in \mathcal{W}, \mathbf{x} \in \mathcal{S}\}$

222 3: $\theta_g^{(0)}, \theta_a^{(0)}, \lambda_g, \lambda_a, \Delta_g, \Delta_a, T_g, T_a \leftarrow \mathcal{M}(b)$

223 4: **for** $t = 1$ to T_g **do**

224 5: $\mathbf{x} \leftarrow \text{Sample}(\mathcal{X}, \mathcal{V})$

225 6: $\theta_g^{(t+1)} \leftarrow \theta_g^{(t)} - \gamma_g \nabla \mathcal{L}_{\text{geometry}}(f_g(\mathbf{x}; \theta_g^{(t)}, \lambda_g))$

226 7: **end for**

227 8: $\hat{\theta}_g \leftarrow \text{Quantization}(\theta_g^{(T_g)}, \Delta_g)$

228 9: $\mathcal{O} \leftarrow \{p \mid f_g(\mathbf{x}; \hat{\theta}_g), \mathbf{x} \in \mathcal{V}\}$

229 10: $\tau \leftarrow \text{AdaptiveThreshold}(\mathcal{O})$

230 11: $\hat{\mathcal{X}} \leftarrow \{\mathbf{x} \mid f_g(\mathbf{x}; \hat{\theta}_g) > \tau, \mathbf{x} \in \mathcal{V}\}$

231 12: **for** $t = 0$ to T_a **do**

232 13: $\mathbf{x} \leftarrow \text{Sample}(\hat{\mathcal{X}})$

233 14: $\theta_a^{(t+1)} \leftarrow \theta_a^{(t)} - \gamma_a \nabla \mathcal{L}_{\text{attribute}}(f_a(\mathbf{x}; \theta_a^{(t)}, \lambda_a))$

234 15: **end for**

235 16: $\hat{\theta}_a \leftarrow \text{Quantization}(\theta_a^{(T_a)}, \Delta_a)$

236 17: $\hat{\theta}_g, \hat{\theta}_a \leftarrow \text{EntropyEncode}(\hat{\theta}_g, \hat{\theta}_a)$

238
239
240 Considering that we train our INRs on GPU and that a single-frame point cloud contains millions
241 of points, the $O(|\mathcal{V}|)$ memory overhead is non-negligible. We split single coordinate \mathbf{x} into two
242 components, which are represented as:

243
$$\mathbf{x} = \mathbf{w} + \mathbf{w}' = \lfloor \mathbf{x} \cdot 2^M \rfloor / 2^M + \mathbf{w}', \quad (12)$$

244
$$\mathbf{w}' = (x, y, z) \in \mathcal{S}' = \left\{ \left(\frac{k_x}{2^N}, \frac{k_y}{2^N}, \frac{k_z}{2^N} \right) \mid k_x, k_y, k_z \in \mathbb{Z}, 0 \leq k_x, k_y, k_z < 2^{N-M} \right\}. \quad (13)$$

245 where \mathbf{w} denotes the coordinates of the non-empty cube containing \mathbf{x} , and \mathbf{w}' represents the relative
246 position of \mathbf{x} within \mathbf{w} . Therefore, the process of sampling a batch of \mathbf{x} can be expressed as:

247
$$\mathbf{x} = \mathbf{w} + \mathbf{w}' = \hat{\alpha} \cdot U(\mathcal{X}) \oplus (1 - \hat{\alpha}) \cdot (U(\mathcal{W}) + U(\mathcal{S}')). \quad (14)$$

248 By decomposing \mathbf{x} and sampling separately from \mathcal{W} and \mathcal{S}' , we avoid explicitly storing \mathcal{V} . In
249 practice, it suffices to generate a set of random coordinates from \mathcal{W} and \mathcal{S}' independently to compute
250 \mathbf{x} . This reduces the storage overhead to $O(|\mathcal{W}|)$. Considering that $O(|\mathcal{W}|)$ typically corresponds to
251 only 0.02% of $O(|\mathcal{V}|)$, this optimization is highly significant.

252

3.3 DYNAMIC THRESHOLD

253 In Eq. 4, we need a threshold τ to classify points as either occupied or unoccupied. A static threshold
254 is unlikely to perform robustly across diverse point clouds. So we introduce a dynamic thresholding
255 mechanism to improve geometric reconstruction. We first define the set of all predicted occupancy
256 probabilities within our sampling space \mathcal{V} as \mathcal{O} :

257
$$\mathcal{O} = \{p \mid p = f_g(\mathbf{x}), \mathbf{x} \in \mathcal{V}\}. \quad (15)$$

258 We use D1 PSNR as the metric for geometry quality, denoted by $\mathcal{D}(\mathcal{O}, \tau)$. Empirically, we observed
259 that this function is typically unimodal with respect to τ in Appendix B. Therefore, a golden section
260 search algorithm can be employed to find the optimal τ that maximizes $\mathcal{D}(\mathcal{O}, \tau)$.

261

3.4 DYNAMIC RATE CONTROL

262 In this section, we describe how `PICO` achieves dynamic rate control from coarse-grained to fine-
263 grained compression. The approach consists of three components: adaptive model selection, regu-
264 larized training, and quantization with entropy coding.

270 **Algorithm 2** `PICO_DeCompression`

271 **Input:** bitstream $\tilde{\theta}_g, \tilde{\theta}_a$, space \mathcal{V}
 272 **Parameter:** threshold τ , model information f_g, f_a
 273 **Output:** decompressed point cloud $\hat{\mathcal{P}} = \{\hat{\mathcal{X}}, \hat{\mathcal{A}}\}$
 274
 275 1: $\hat{\theta}_g, \hat{\theta}_a \leftarrow \text{EntropyDecode}(\tilde{\theta}_g, \tilde{\theta}_a)$
 276 2: $\hat{\mathcal{X}} \leftarrow \{\mathbf{x} \mid f_g(\mathbf{x}; \hat{\theta}_g) > \tau, \mathbf{x} \in \mathcal{V}\}$
 277 3: $\hat{\mathcal{A}} \leftarrow \{f_a(\mathbf{x}; \hat{\theta}_a) \mid \mathbf{x} \in \hat{\mathcal{X}}\}$
 278 4: $\hat{\mathcal{P}} \leftarrow \{\hat{\mathcal{X}}, \hat{\mathcal{A}}\}$

280

281 **Adaptive Model Selection.** Previous works often employ a single network architecture to compress
 282 point clouds of varying sizes. This approach can lead to significant performance degradation when
 283 the target bitrate changes. We observe that while small-parameter models struggle to compete with
 284 large-parameter models in compression quality at low compression rates, they gradually surpass
 285 larger models as the compression rate increases. This observation naturally led us to design a model
 286 dictionary for selecting the optimal model f based on the target compression rate.

287 To construct this dictionary, we compress the same point cloud at different rates using models with
 288 various parameter counts. From their rate-distortion (RD) curves, we derive a Pareto frontier to
 289 define our model dictionary \mathcal{M} . In our experiments, we found that a model dictionary obtained from
 290 a single experiment can be broadly applied to a wide range of point clouds. For excessively large
 291 point clouds, we divide them into smaller blocks for processing. This ensures the high efficiency
 292 and practicality of our method for real-world deployment.

293 **Regularized Training.** We employ different objective functions for the two compression stages.
 294 For geometry compression, we use an α -modulated focal loss Lin et al. (2017), which effectively
 295 addresses the class imbalance problem through a weighting mechanism. The balancing factor fur-
 296 ther corrects for the label imbalance between empty and occupied voxels caused by sampling bias.
 297 For attribute compression, we use a per-voxel mean squared error (MSE) loss, which is calculated
 298 only on occupied voxels to ensure perceptually accurate attributes. Furthermore, we add an ℓ_1 regu-
 299 larization term during training to control the sparsity and distribution of the model parameters. The
 300 complete loss functions for geometry and attribute compression are as follows:

$$\mathcal{L}_{\text{geometry}} = \mathcal{L}_{\text{focal}} + \mathcal{L}_{\text{reg}} = \mathbb{E}_{p \sim \mathcal{O}}[-\alpha_t(1 - p_t)^\gamma \log(p_t)] + \lambda_g \|\Theta_g\|_1, \quad (16)$$

$$\mathcal{L}_{\text{attribute}} = \mathcal{L}_{\text{MSE}} + \mathcal{L}_{\text{reg}} = \mathbb{E}_{\mathbf{x} \sim \hat{\mathcal{X}}}[\|\tilde{\mathcal{A}}(\mathbf{x}_i) - f_a(\mathbf{x}_i)\|_2^2] + \lambda_a \|\Theta_a\|_1. \quad (17)$$

304 **Quantization & Entropy Coding.** Once the model parameters are obtained, they are quantized with
 305 step sizes Δ_g and Δ_a for geometry and attributes, respectively, and then entropy-coded to generate
 306 the final bitstream. Our framework utilizes DeepCABAC (Wiedemann et al., 2019), a specialized
 307 context-adaptive binary arithmetic coder designed for deep neural networks.

308 Suppose the model has K parameters, denoted by θ , with ℓ_1 regularization strength λ and a quanti-
 309 zation step size Δ . Then the average code length of the final bitstream of θ after entropy coding can
 310 be expressed as:

$$H(\theta) \approx K \log \frac{2e}{\lambda \Delta}. \quad (18)$$

311 We provide a derivation of this result in the Appendix C.

315 3.5 DYNAMIC POINT CLOUD COMPRESSION

317 We generalize the `PICO` framework to dynamic PCC, modeling the point cloud frame sequence
 318 as a function over a 4D spatio-temporal domain. By exploiting temporal redundancy, we augment
 319 the spatial coordinates (x, y, z) with a time coordinate t . Consequently, our INRs become map-
 320 pings which take a spatio-temporal coordinate (x, y, z, t) as input and outputs a tuple containing the
 321 geometric occupancy $p \in \mathbb{R}$ or the attribute $\mathbf{c} \in \mathbb{R}^3$. The `PICO-Dynamic` can be expressed as:

$$f_g : (x, y, z, t) \mapsto p \in [0, 1], \quad \hat{\mathcal{X}}^{(t)} = \{\mathbf{x} \mid f_g(\mathbf{x}) > \tau_0^{(t)}, \mathbf{x} \in \mathcal{V}^{(t)}\}, \quad (19)$$

$$f_a : (x, y, z, t) \mapsto \mathbf{c} \in [0, 1]^3, \quad \hat{\mathcal{A}}^{(t)} = \{f_a(\mathbf{x}) \mid \mathbf{x} \in \hat{\mathcal{X}}^{(t)}\}. \quad (20)$$

324 3.6 LEARNABLE ACTIVATION FUNCTION NETWORK
325

326 We propose an INR backbone LeAFNet tai-
327 loled for PCC in PICO. LeAFNet is designed
328 to enhance the fitting of implicit functions for
329 both geometric and attribute compression. The
330 network structure is illustrated in Figure 1.

331 **Positional Encoding.** LeAFNet processes 3D
332 voxel coordinates through NeRF (Mildenhall
333 et al., 2021)-style positional encoding to ad-
334 dress the limited representational capacity of
335 raw coordinate inputs. Specifically, each co-
336 ordinate \mathbf{x} is mapped to a higher-dimensional
337 vector using a series of sinusoidal functions
338 before being input to the network. This trans-
339 formation allows the network to capture fine-
340 grained geometric details and subtle attribute
341 variations that would otherwise be lost in the
342 raw representation, which is critical for high-
343 fidelity reconstruction. The positional encoding
344 Γ can be expressed as:

$$\Gamma(\mathbf{x}; L) = (\mathbf{x}, \sin(2^0 \pi \mathbf{x}), \cos(2^0 \pi \mathbf{x}), \dots, \sin(2^{L-1} \pi \mathbf{x}), \cos(2^{L-1} \pi \mathbf{x})). \quad (21)$$

346 **Learnable Activation Layer.** The design of LeAFNet aims to integrate learnable activation func-
347 tions to enhance implicit function approximation while maintaining parameter efficiency. The first
348 part of LeAFNet is several fully connected layers, which reduce the dimensionality of the positional
349 encoded input for subsequent processing. The core of LeafNet is the learnable activation function
350 layer. Here, we adopt an approach similar to KAN, using the same structure as the KAN_Layer.
351 The difference is that we replace the low-throughput B-spline functions with efficient radial basis
352 functions (Li, 2024) (RBFs). The learnable activation function $\phi(x)$ is therefore defined as:

$$\phi(x) = w_b \text{silu}(x) + w_s \sum_{i=1}^N \exp\left(-\frac{\|x - c_i\|^2}{h^2}\right), \quad (22)$$

356 where c_i and h are hyperparameters that determine the shape of the activation function, and w_b
357 and w_s are learnable linear layer weights. This design significantly accelerates both forward and
358 backward passes, enabling faster computation and more stable training compared to KAN.
359

360 4 EXPERIMENTS
361362 4.1 EXPERIMENTAL SETUPS
363

364 **Baselines.** We selected six state-of-the-art methods as our baselines for comparison: G-PCC
365 (Graziosi et al., 2020), V-PCC (Graziosi et al., 2020), NeRC (Ruan et al., 2024b), NeRI (Xue et al.,
366 2024), SparsePCGC (Wang et al., 2022), and Unicorn (Wang et al., 2024).

367 **Datasets** We conducted experiments on three widely recognized point cloud datasets: 8iVFB (d’Eon
368 et al., 2019), KITTI (Geiger et al., 2012), and ScanNet (Dai et al., 2017b).

370 **Metrics.** In our experiments, we use the point-to-point error peak signal-to-noise ratio (D1 PSNR)
371 (dB) to measure geometric distortion across all three datasets. For attribute distortion, we use
372 dataset-specific metrics: **PCQM** (Meynet et al., 2020) ($\times 10^{-3}$) for the 8iVFB dataset, **R-PSNR**
373 (dB) for KITTI, and **Y-PSNR** (dB) for ScanNet. To quantify the rate-distortion (RD) performance
374 gains of different methods, we utilize the Bjontegaard delta metrics (Bjontegaard, 2001).

375 **Implementation Details.** We set the coarse-grained voxel resolution to $M = 5$, partitioning the
376 space into $2^5 \times 2^5 \times 2^5$ cubes, each containing $32 \times 32 \times 32$ voxels. The network is optimized
377 by Adam (Kingma & Ba, 2014), with an initial learning rate of 1×10^{-3} , which is decayed by a
378 factor of 0.1 upon reaching a performance plateau. A batch size of 32,768 is used for training. For

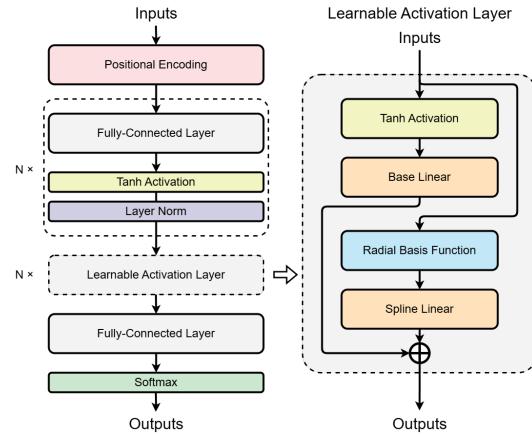


Figure 1: Model Architecture of LeAFNet.

378	379	Method	KITTI	ScanNet	8iVFB				
					380	longdress	loot	redandblack	soldier
G-PCC (octree)	-34.22 / 5.14	-78.65 / 6.51	-69.19 / 7.33	-71.78 / 7.84	-69.25 / 6.48	-60.76 / 6.69	-67.75 / 7.09		
G-PCC (trisoup)	-28.25 / 4.05	-71.24 / 5.29	-49.34 / 3.96	-53.33 / 4.36	-47.20 / 3.33	-42.75 / 4.15	-48.16 / 3.95		
V-PCC	-15.64 / 2.88	-56.42 / 3.67	-35.42 / 2.64	-44.66 / 3.64	-53.27 / 4.42	-45.51 / 4.20	-44.72 / 3.73		
NeRC	-18.57 / 3.65	-58.23 / 4.05	-33.13 / 2.80	-45.53 / 4.56	-48.12 / 4.92	-43.55 / 4.36	-42.58 / 4.16		
NeRI	-20.14 / 3.28	-54.78 / 3.02	-31.28 / 2.64	-49.89 / 6.54	-48.65 / 6.43	-45.47 / 5.91	-43.89 / 5.28		
SparsePCGC	-8.94 / 1.13	-51.98 / 3.81	-17.23 / 0.67	-22.54 / 1.25	-26.51 / 1.03	-18.28 / 1.29	-21.14 / 1.06		
Unicorn	0.66 / -0.16	2.83 / -0.25	13.24 / -1.22	11.15 / -0.43	10.98 / -0.61	9.27 / -0.98	11.16 / -0.81		

386	387	Method	KITTI	ScanNet	8iVFB				
					388	longdress	loot	redandblack	soldier
G-PCC (octree)	-13.48 / 1.24	-25.61 / 2.32	-40.94 / 3.89	-74.72 / 6.86	-71.68 / 5.86	-65.87 / 7.69	-63.30 / 6.08		
G-PCC (trisoup)	-9.43 / 0.81	-21.20 / 1.28	-19.88 / 0.88	-47.68 / 1.43	-40.30 / 1.73	-57.49 / 2.37	-41.34 / 1.60		
V-PCC	-4.17 / 0.36	-8.63 / 0.55	-8.19 / 0.68	-29.53 / 0.22	-46.48 / 1.15	-26.18 / 1.08	-27.60 / 0.78		
NeRC	-5.20 / 0.54	-7.09 / 0.39	-17.19 / 1.35	-49.41 / 2.51	-51.77 / 2.51	-43.68 / 3.06	-40.51 / 2.36		
NeRI	-9.20 / 1.25	-10.10 / 0.51	-23.48 / 1.88	-48.41 / 3.78	-46.37 / 3.42	-28.74 / 5.64	-36.75 / 3.68		
SparsePCGC	-0.53 / 0.05	-15.52 / 0.94	-12.92 / 0.88	-18.49 / 1.56	-17.83 / 1.12	-17.04 / 1.92	-16.57 / 1.37		
Unicorn	19.64 / 0.30	2.86 / -0.09	10.32 / -2.31	8.95 / 0.14	6.13 / -0.88	7.48 / -0.23	8.22 / -0.89		

Table 1: We evaluated **PICO**’s compression performance against seven baselines on three datasets. The top part of the table shows the results for geometry-only compression, while the bottom part displays the results for joint geometry and attribute compression. We report the Bjontegaard delta (BD) gains of **PICO** relative to baseline methods. The number before the slash indicates the BD-Rate (%), where a lower value is better. The number after the slash represents the corresponding metric, where a higher value is better. RD curve visualization can be found in Appendix A.

geometry compression, we set the reweighted sampling coefficient to $\alpha = 0.5$ and the focal loss modulation coefficient to $\gamma = 2$. The quantized parameters from both models are then losslessly compressed using DeepCABAC (Wiedemann et al., 2019).

4.2 POINT CLOUD COMPRESSION

Static Point Cloud. To evaluate **PICO**’s performance, we selected a single frame from each of the four point cloud sequences in the 8iVFB dataset, and used individual point clouds from the KITTI and ScanNet datasets. The results are shown in Table 1. **PICO** achieves substantial improvements over conventional MPEG standards (G-PCC and V-PCC), delivering higher compression efficiency and reconstruction quality. It also surpasses existing INR-based approaches (NeRC and NeRI) and demonstrates a clear advantage over SparsePCGC. Nevertheless, there remains a small performance gap compared to the current state-of-the-art, Unicorn.

PICO’s improvements are especially notable in geometry compression, thanks to its continuous occupancy representation, which fits well with the nature of INRs. Attribute compression is more challenging, mainly because geometric errors accumulate and sharp attribute changes are difficult to capture with the smooth mappings of INR.

Dynamic Point Cloud. To evaluate **PICO**-dynamic’s performance, we selected the first 30 frames of the 8iVFB_longdress sequence. We benchmarked **PICO** against state-of-the-art methods, including V-PCC, NeRC, and Unicorn. The experimental results in Table 2 clearly demonstrate that **PICO** exhibits superior compression performance on dynamic data. **PICO** consistently outperforms both V-PCC and NeRC. A noteworthy finding is that **PICO** successfully surpasses Unicorn in geometry compression, highlighting the strength of our spatio-temporal geometry representation.

Method	8iVFB.longdress		
	Geo. Only	BD-BR (%)	BD-PSNR (dB)
V-PCC	-59.51	5.30	
NeRC	-46.25	4.38	
Unicorn	-3.68	0.24	

Method	8iVFB.longdress		
	Geo. & Attr.	BD-BR (%)	BD-PCQM ($\times 10^{-3}$)
V-PCC	-23.98	1.25	
NeRC	-26.81	2.14	
Unicorn	6.22	-0.51	

Table 2: **PICO** vs. baselines on the first 30 frames of 8iVFB_longdress. Top: geometry compression. Bottom: joint compression.

432 4.3 IMPACT OF LEAFNET
433

Parameters			8iVFB_loot	
<i>d</i>	<i>L</i>	α	BD-BR (%)	BD-PSNR (dB)
24	64	0.5	25.15	0.61
48	64	0.5	-11.16	-0.11
36	48	0.5	-5.36	0.89
36	32	0.5	-25.64	1.36
36	16	0.5	-34.21	2.68
36	0	0.5	-41.63	3.38
36	64	0.25	-39.09	2.93
36	64	δ	-61.78	5.63

446 Table 3: Ablation study of LeAFNet’s hyperpa-
447 rameters. Default $\{d, L, \alpha\} = \{36, 64, 0.5\}$

449 the positive label ratio α from its original, highly imbalanced value δ to 0.5 leads to a substantial
450 gain in compression quality, as this resampling strategy mitigates training instability and facilitates
451 learning a more robust occupancy probability distribution. The hidden layer dimension d also plays
452 an important role: larger values of d increase the bitrate but simultaneously improve performance,
453 highlighting the effective scalability of the LeAFNet architecture.

454 4.3.1 ADAPTIVE MODEL PARAMETER SELECTION
455

456 To provide a more intuitive understanding of
457 the model dictionary \mathcal{M} , we constructed a
458 model dictionary on the 8iVFB_longdress,
459 as illustrated in Figure 2. We varied the hid-
460 den layer dimensions $d \in \{16, 32, 48, 64, 80, 96\}$
461 and network depths (number of layers) \in
462 $\{2, 3, 4, 5, 6, 7\}$, and assigned different train-
463 ing hyperparameters to each configuration to
464 achieve a balance between performance and
465 compression efficiency. Our experiments con-
466 firm the trends observed in earlier analyses: at
467 lower bitrates, smaller models tend to perform
468 better, as excessive quantization steps Δ and
469 strong ℓ_1 regularization λ can degrade the rep-
470 resentational capacity of larger models; con-
471 versely, at higher bitrates, larger models dom-
472 inate, with their increased parameter capacity
473 allowing them to better capture fine-grained point cloud details.

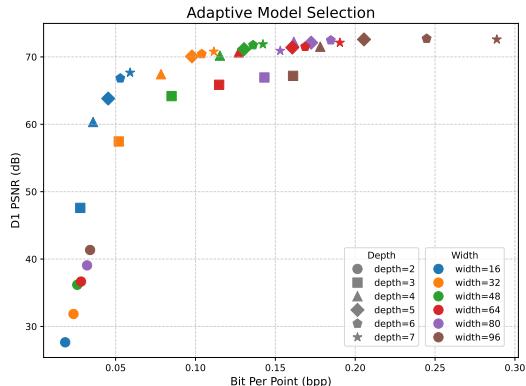
474 5 CONCLUSION
475

476 In this work, we introduced PICO, a INR-based PCC framework that decouples geometry and at-
477 tribute compression. By reformulating PCC as a neural network compression, PICO achieves flex-
478 ible control over bitrate and reconstruction quality. We further proposed LeAFNet, a lightweight
479 INR backbone with learnable activation functions, positional encoding and radial basis functions,
480 which effectively capture high-frequency point cloud details with fewer parameters.

481 Extensive experiments on 8iVFB, KITTI, and ScanNet datasets demonstrate that PICO consistently
482 outperforms traditional MPEG standards and existing PCC methods, achieving substantial gains in
483 both geometry and joint compression metrics. Overall, PICO represents a significant step toward
484 next-generation point cloud compression, providing a flexible, high-performance, and deployable
485 solution that bridges INR and practical compression needs.

We conducted a comprehensive ablation study to investigate the impact of LeAFNet’s hyperparameters—hidden layer dimension d , positional encoding dimension L , and sampling ratio α —on compression performance. For this analysis, we established a default configuration with $d = 36$, $L = 64$, and $\alpha = 0.5$, and then varied each hyperparameter to evaluate its individual and combined effects on PICO. The results are summarized in Table 3.

Our analysis indicates that both the positional encoding and the sampling strategy are crucial for enhancing LeAFNet’s performance. Increasing the positional encoding dimension L improves compression by enabling the model to capture finer spatial details. Likewise, adjusting

Figure 2: The 8iVFB_longdress model dic-
tionary, also applicable to other point clouds.

486 REFERENCES
487

488 Gisle Bjontegaard. Calculation of average psnr differences between rd-curves. *ITU SG16 Doc.*
489 *VCEG-M33*, 2001.

490 Yaodong Cui, Ren Chen, Wenbo Chu, Long Chen, Daxin Tian, Ying Li, and Dongpu Cao. Deep
491 learning for image and point cloud fusion in autonomous driving: A review. *IEEE Transactions*
492 *on Intelligent Transportation Systems*, 23(2):722–739, 2021.

493 Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias
494 Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the*
495 *IEEE conference on computer vision and pattern recognition*, pp. 5828–5839, 2017a.

496 Angela Dai, Matthias Nießner, Michael Zollöfer, Shahram Izadi, and Christian Theobalt. Bundle-
497 fusion: Real-time globally consistent 3d reconstruction using on-the-fly surface re-integration.
498 *ACM Transactions on Graphics 2017 (TOG)*, 2017b.

499 Ricardo L De Queiroz and Philip A Chou. Compression of 3d point clouds using a region-adaptive
500 hierarchical transform. *IEEE Transactions on Image Processing*, 25(8):3947–3956, 2016.

501 Emilien Dupont, Adam Goliński, Milad Alizadeh, Yee Whye Teh, and Arnaud Doucet. Coin: Com-
502 pression with implicit neural representations. *arXiv preprint arXiv:2103.03123*, 2021.

503 Emilien Dupont, Hrushikesh Loya, Milad Alizadeh, Adam Goliński, Yee Whye Teh, and Arnaud
504 Doucet. Coin++: Neural compression across modalities. *arXiv preprint arXiv:2201.12904*, 2022.

505 Eugene d’Eon, Bob Harrison, Taos Myers, and Philip A Chou. Jpeg pleno database: 8i voxelized
506 full bodies (8ivfb v2)-a dynamic voxelized point cloud dataset, 2019.

507 Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti
508 vision benchmark suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)*,
509 2012.

510 Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The
511 kitti dataset. *The international journal of robotics research*, 32(11):1231–1237, 2013.

512 Danillo Graziosi, Ohji Nakagami, Shinroku Kuma, Alexandre Zaghetto, Teruhiko Suzuki, and Ali
513 Tabatabai. An overview of ongoing point cloud compression standardization activities: Video-
514 based (v-pcc) and geometry-based (g-pcc). *APSIPA Transactions on Signal and Information*
515 *Processing*, 9:e13, 2020.

516 Tianxin Huang and Yong Liu. 3d point cloud geometry compression on deep learning. In *Proceed-
517 ings of the 27th ACM international conference on multimedia*, pp. 890–898, 2019.

518 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*
519 *arXiv:1412.6980*, 2014.

520 Ying Li, Lingfei Ma, Zilong Zhong, Fei Liu, Michael A Chapman, Dongpu Cao, and Jonathan Li.
521 Deep learning for lidar point clouds in autonomous driving: A review. *IEEE Transactions on*
522 *Neural Networks and Learning Systems*, 32(8):3412–3432, 2020.

523 Ziyao Li. Kolmogorov-arnold networks are radial basis function networks. *arXiv preprint*
524 *arXiv:2405.06721*, 2024.

525 Sohee Lim, Minwoo Shin, and Joonki Paik. Point cloud generation using deep adversarial local
526 features for augmented and mixed reality contents. *IEEE Transactions on Consumer Electronics*,
527 68(1):69–76, 2022.

528 Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense
529 object detection. In *Proceedings of the IEEE international conference on computer vision*, pp.
530 2980–2988, 2017.

531 Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljacic,
532 Thomas Y. Hou, and Max Tegmark. KAN: Kolmogorov–arnold networks. In *The Thirteenth*
533 *International Conference on Learning Representations*, 2025.

540 Gabriel Meynet, Yana Nehmé, Julie Digne, and Guillaume Lavoué. Pcq: A full-reference quality
 541 metric for colored 3d point clouds. In *2020 Twelfth International Conference on Quality of*
 542 *Multimedia Experience (QoMEX)*, pp. 1–6. IEEE, 2020.

543

544 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and
 545 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications*
 546 *of the ACM*, 65(1):99–106, 2021.

547

548 Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets
 549 for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision*
 550 *and pattern recognition*, pp. 652–660, 2017.

551

552 Zekun Qi, Runpei Dong, Shaochen Zhang, Haoran Geng, Chunrui Han, Zheng Ge, Li Yi, and
 553 Kaisheng Ma. Shapellm: Universal 3d object understanding for embodied interaction. In *Europ-
 554 ean Conference on Computer Vision*, pp. 214–238. Springer, 2024.

555

556 Maurice Quach, Giuseppe Valenzise, and Frederic Dufaux. Improved deep point cloud geom-
 557 etry compression. In *2020 IEEE 22nd International Workshop on Multimedia Signal Processing
 (MMSP)*, pp. 1–6. IEEE, 2020.

558

559 Maurice Quach, Jiahao Pang, Dong Tian, Giuseppe Valenzise, and Frédéric Dufaux. Survey on deep
 560 learning-based point cloud compression. *Frontiers in Signal Processing*, 2:846972, 2022.

561

562 Zizheng Que, Guo Lu, and Dong Xu. Voxelcontext-net: An octree based framework for point
 563 cloud compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern
 Recognition*, pp. 6042–6051, 2021.

564

565 Thinal Raj, Fazida Hanim Hashim, Aqilah Baseri Huddin, Mohd Faisal Ibrahim, and Aini Hussain.
 566 A survey on lidar scanning mechanisms. *Electronics*, 9(5):741, 2020.

567

568 Sameera Ramasinghe and Simon Lucey. Beyond periodicity: Towards a unifying framework for acti-
 569 vations in coordinate-mlps. In *European Conference on Computer Vision*, pp. 142–158. Springer,
 2022.

570

571 Hongning Ruan, Yulin Shao, Qianqian Yang, Liang Zhao, and Dusit Niyato. Point cloud compres-
 572 sion with implicit neural representations: A unified framework. In *2024 IEEE/CIC International
 573 Conference on Communications in China (ICCC)*, pp. 1709–1714. IEEE, 2024a.

574

575 Hongning Ruan, Yulin Shao, Qianqian Yang, Liang Zhao, Zhaoyang Zhang, and Dusit Niyato.
 576 Implicit neural compression of point clouds. *arXiv preprint arXiv:2412.10433*, 2024b.

577

578 Vishwanath Saragadam, Daniel LeJeune, Jasper Tan, Guha Balakrishnan, Ashok Veeraraghavan,
 579 and Richard G Baraniuk. Wire: Wavelet implicit neural representations. In *Proceedings of the
 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18507–18516, 2023.

580

581 Sebastian Schwarz, Marius Preda, Vittorio Baroncini, Madhukar Budagavi, Pablo Cesar, Philip A
 582 Chou, Robert A Cohen, Maja Krivokuća, Sébastien Lasserre, Zhu Li, et al. Emerging mpeg
 583 standards for point cloud compression. *IEEE Journal on Emerging and Selected Topics in Circuits
 and Systems*, 9(1):133–148, 2018.

584

585 Xihua Sheng, Li Li, Dong Liu, Zhiwei Xiong, Zhu Li, and Feng Wu. Deep-pcac: An end-to-end
 586 deep lossy compression framework for point cloud attributes. *IEEE Transactions on Multimedia*,
 587 24:2617–2632, 2021.

588

589 Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Im-
 590 plicit neural representations with periodic activation functions. *Advances in neural information
 591 processing systems*, 33:7462–7473, 2020.

592

593 Yannick Strümpler, Janis Postels, Ren Yang, Luc Van Gool, and Federico Tombari. Implicit neural
 594 representations for image compression. In *European Conference on Computer Vision*, pp. 74–91.
 Springer, 2022.

594 Gary J Sullivan, Jens-Rainer Ohm, Woo-Jin Han, and Thomas Wiegand. Overview of the high
 595 efficiency video coding (hevc) standard. *IEEE Transactions on circuits and systems for video*
 596 *technology*, 22(12):1649–1668, 2012.

597

598 Jianqiang Wang, Dandan Ding, Zhu Li, and Zhan Ma. Multiscale point cloud geometry compression.
 599 In *2021 Data Compression Conference (DCC)*, pp. 73–82. IEEE, 2021a.

600 Jianqiang Wang, Hao Zhu, Haojie Liu, and Zhan Ma. Lossy point cloud geometry compression via
 601 end-to-end learning. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(12):
 602 4909–4923, 2021b.

603

604 Jianqiang Wang, Dandan Ding, Zhu Li, Xiaoxing Feng, Chunlong Cao, and Zhan Ma. Sparse
 605 tensor-based multiscale representation for point cloud geometry compression. *IEEE Transactions*
 606 *on Pattern Analysis and Machine Intelligence*, 45(7):9055–9071, 2022.

607 Jianqiang Wang, Ruixiang Xue, Jiaxin Li, Dandan Ding, Yi Lin, and Zhan Ma. A versatile point
 608 cloud compressor using universal multiscale conditional coding–part i: Geometry. *IEEE transac-*
 609 *tions on pattern analysis and machine intelligence*, 2024.

610 Zeyu Wang, Cuong Nguyen, Paul Asente, and Julie Dorsey. Pointshopar: Supporting environmental
 611 design prototyping using point cloud in augmented reality. In *Proceedings of the 2023 CHI*
 612 *Conference on Human Factors in Computing Systems*, pp. 1–15, 2023.

613

614 Simon Wiedemann, Heiner Kirchhoffer, Stefan Matlage, Paul Haase, Arturo Marban, Talmaj
 615 Marinc, David Neumann, Ahmed Osman, Detlev Marpe, Heiko Schwarz, et al. Deepcabac:
 616 Context-adaptive binary arithmetic coding for deep neural network compression. *arXiv preprint*
 617 *arXiv:1905.08318*, 2019.

618 Ruixiang Xue, Jiaxin Li, Tong Chen, Dandan Ding, Xun Cao, and Zhan Ma. Neri: Implicit neural
 619 representation of lidar point cloud using range image sequence. In *ICASSP 2024-2024 IEEE*
 620 *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8020–8024.
 621 IEEE, 2024.

622

623 Junzhe Zhang, Gexin Liu, Junteng Zhang, Dandan Ding, and Zhan Ma. Deppcc: Learned lossy
 624 point cloud compression. *IEEE Transactions on Emerging Topics in Computational Intelligence*,
 625 2024.

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

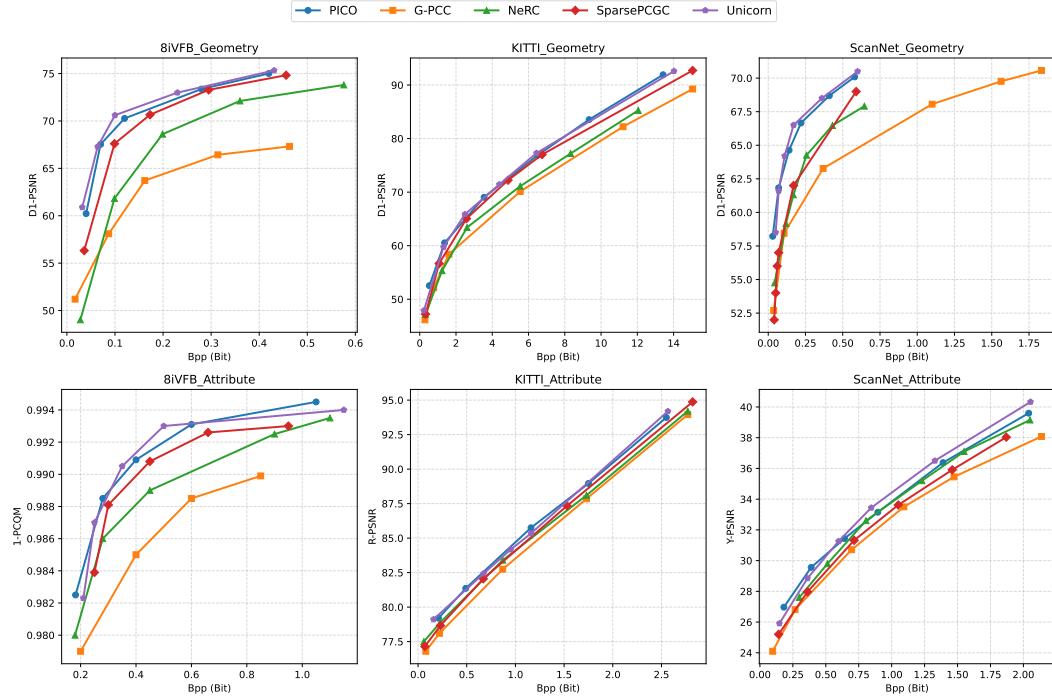
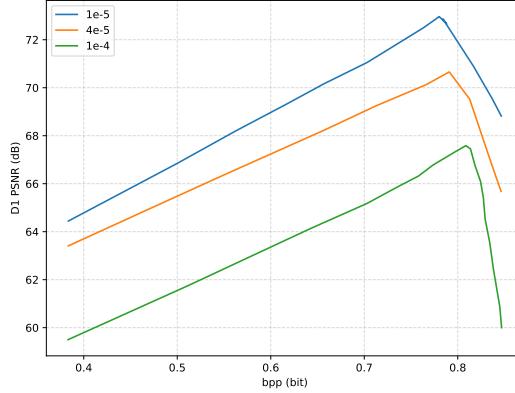
645

646

647

648
649 A RD CURVE VISUALIZATION

650
651 In Figure 3, we visualize the RD curves on three static point clouds, including the five methods:
652 PICO (Ours), G-PCC (Graziosi et al., 2020), NeRC (Ruan et al., 2024b), SparsePCGC (Wang et al.,
653 2022), and Unicorn (Wang et al., 2024). It can be found that PICO is basically on par with the state-
654 of-the-art Unicorn, showing a significant performance improvement compared to the other methods.
655

666
667 Figure 3: RD Curves of five methods on three different point clouds.
668669
670 B UNIMODALITY OF $D(\mathcal{O}, \tau)$ 696
697 Figure 4: Unimodality of $D(\mathcal{O}, \tau)$
698

699 In Figure 4, we visualize a set of $D(\mathcal{O}, \tau)$ curves. For the same model, the D curves under different
700 regularization strengths all exhibit unimodality with respect to τ . Moreover, it can be observed that
701 by adjusting τ , we can achieve up to an 8 dB gain, demonstrating the effectiveness of our dynamic
thresholding method.
702

702 **C DERIVATION OF $H(\theta) \approx K \log \frac{2e}{\lambda\Delta}$**

704 **Theorem 1** (Average coding length of quantized ℓ_1 -regularized parameters). *Suppose a model has*
 705 *K parameters $\theta = (\theta_1, \dots, \theta_K)$, each subject to ℓ_1 regularization with strength λ , and the param-*
 706 *eters are quantized with step size Δ . Then the average code length of the final bitstream of θ after*
 707 *entropy coding is approximately*

$$708 \quad 709 \quad 710 \quad 711 \quad 712 \quad 713 \quad 714 \quad 715 \quad 716 \quad 717 \quad 718 \quad 719 \quad 720 \quad 721 \quad 722 \quad 723 \quad 724 \quad 725 \quad 726 \quad 727 \quad 728 \quad 729 \quad 730 \quad 731 \quad 732 \quad 733 \quad 734 \quad 735 \quad 736 \quad 737 \quad 738 \quad 739 \quad 740 \quad 741 \quad 742 \quad 743 \quad 744 \quad 745 \quad 746 \quad 747 \quad 748 \quad 749 \quad 750 \quad 751 \quad 752 \quad 753 \quad 754 \quad 755$$

$$H(\theta) \approx K \log \frac{2e}{\lambda\Delta}.$$

Proof. We provide a detailed derivation of Eq.18 as follows.

Consider a single model parameter θ_i under ℓ_1 regularization. Its prior distribution is the Laplace distribution, with probability density function (PDF):

$$p(\theta_i | \lambda) = \frac{\lambda}{2} \exp(-\lambda|\theta_i|),$$

where λ is the regularization strength. This distribution forms the basis of our entropy calculation.

The differential entropy of a continuous random variable X is defined as

$$H(X) = - \int_{-\infty}^{\infty} p(x) \log p(x) dx.$$

Substituting the Laplace PDF, we have

$$\begin{aligned} H(\theta_i) &= - \int_{-\infty}^{\infty} \frac{\lambda}{2} e^{-\lambda|\theta_i|} \log \left(\frac{\lambda}{2} e^{-\lambda|\theta_i|} \right) d\theta_i \\ &= - \int_{-\infty}^{\infty} \frac{\lambda}{2} e^{-\lambda|\theta_i|} \left[\log \left(\frac{\lambda}{2} \right) - \lambda|\theta_i| \right] d\theta_i \\ &= - \left[\log \left(\frac{\lambda}{2} \right) \int_{-\infty}^{\infty} \frac{\lambda}{2} e^{-\lambda|\theta_i|} d\theta_i - \lambda \int_{-\infty}^{\infty} |\theta_i| \frac{\lambda}{2} e^{-\lambda|\theta_i|} d\theta_i \right]. \end{aligned}$$

The first integral evaluates to 1, because it is the integral of the PDF over the whole real line. The second integral is the expected value of $|\theta_i|$ under the Laplace distribution, which is $E[|\theta_i|] = 1/\lambda$. Therefore, we obtain

$$H(\theta_i) = - \left[\log \left(\frac{\lambda}{2} \right) - \lambda \cdot \frac{1}{\lambda} \right] = \log \left(\frac{2e}{\lambda} \right),$$

which is the differential entropy of a single parameter.

If we quantize θ_i with a small step size Δ , the discrete entropy of the quantized variable, $H_\Delta(\theta_i)$, is approximately related to the differential entropy by

$$H_\Delta(\theta_i) \approx H(\theta_i) - \log \Delta = \log \left(\frac{2e}{\lambda\Delta} \right).$$

Considering the K parameters $\theta_1, \dots, \theta_K$ are drawn from the same distribution θ obtained from a single training run, the total entropy of the parameter set is the sum of the individual entropies:

$$H(\theta) = \sum_{i=1}^K H_\Delta(\theta_i) \approx \sum_{i=1}^K \log \left(\frac{2e}{\lambda\Delta} \right) = K \log \left(\frac{2e}{\lambda\Delta} \right).$$

This final expression approximates the total number of nats required to encode all quantized parameters optimally. \square

756 **D BASELINES & DATASETS**
757758 We selected six methods as our baselines for comparison, representing a diverse range of point cloud
759 compression techniques.
760761

- 762 • **G-PCC** (Graziosi et al., 2020) is a geometry-based PCC standard that uses an octree struc-
763 ture and entropy coding to compress voxelized geometry and attributes.
764
- 765 • **V-PCC** (Graziosi et al., 2020) is a video-based PCC standard that projects 3D point clouds
766 onto 2D planes for compression using video codecs.
767
- 768 • **NeRC** (Ruan et al., 2024b) is a PCC framework that uses two separate neural networks to
769 implicitly represent point cloud’s geometry and attributes.
770
- 771 • **NeRI** (Xue et al., 2024) compresses point clouds by projecting 3D frames into 2D range
772 images and encoding them via an implicit neural network.
773
- 774 • **SparsePCGC** (Wang et al., 2022) is a multiscale sparse tensor point cloud geometry com-
775 pression method.
776
- 777 • **Unicorn** (Wang et al., 2024) is a versatile, multiscale conditional coding framework that
778 uses spatial and temporal scale priors to jointly compress point clouds.
779

780 We conducted experiments on three widely recognized point cloud datasets to demonstrate the ver-
781 satility of our method and network.
782783

- 784 • **8iVFB** (d’Eon et al., 2019) is a dynamic voxelized point cloud dataset containing sequences
785 of human subjects captured at high resolution, primarily used for evaluating point cloud
786 compression standards like MPEG.
787
- 788 • **KITTI** (Geiger et al., 2012) is a popular dataset for autonomous driving research, featuring
789 a rich collection of outdoor urban scenes with synchronized data from multiple sensors,
790 including a 3D LiDAR scanner, stereo cameras, and GPS/IMU.
791
- 792 • **ScanNet** (Dai et al., 2017b) is a large-scale RGB-D video dataset of indoor scenes, provid-
793 ing richly annotated 3D reconstructions with instance-level semantic segmentations, used
794 for various 3D scene understanding tasks.
795

796 **E GENERATIVE AI USAGE STATEMENT**
797798 In the preparation of this document, We have used generative AI tools only for grammar checking
799 and language refinement. All content, ideas, and technical material were developed independently,
800 and the AI was not used to generate original content, perform substantive analysis, or contribute to
801 the intellectual work itself.
802