
NeCGS: Neural Compression for 3D Geometry Sets

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 This paper explores the problem of effectively compressing 3D geometry sets
2 containing diverse categories. We make *the first* attempt to tackle this fundamental
3 and challenging problem and propose NeCGS, a neural compression paradigm,
4 which can compress hundreds of detailed and diverse 3D mesh models (~684 MB)
5 by about 900 times (0.76 MB) with high accuracy and preservation of detailed
6 geometric details. Specifically, we first represent each *irregular* mesh model/shape
7 in a *regular* representation that implicitly describes the geometry structure of the
8 model using a 4D regular volume, called *TSDf-Def volume*. Such a regular rep-
9 resentation can not only capture local surfaces more effectively but also facilitate
10 the subsequent process. Then we construct a quantization-aware auto-decoder
11 network architecture to regress these 4D volumes, which can summarize the sim-
12 ilarity of local geometric structures within a model and across different models
13 for redundancy elimination, resulting in more compact representations, including
14 an embedded feature of a smaller size associated with each model and a network
15 parameter set shared by all models. We finally encode the resulting features and
16 network parameters into bitstreams through entropy coding. After decompressing
17 the features and network parameters, we can reconstruct the TSDf-Def volumes,
18 where the 3D surfaces can be extracted through the deformable marching cubes.
19 Extensive experiments and ablation studies demonstrate the significant advantages
20 of our NeCGS over state-of-the-art methods both quantitatively and qualitatively.
21 *We have included the source code in the Supplemental Material.*

22 1 Introduction

23 3D mesh models/shapes are widely used in various fields, such as computer graphics, virtual reality,
24 robotics, and autonomous driving. As geometric data becomes increasingly complex and voluminous,
25 effective compression techniques have become critical for efficient storage and transmission. More-
26 over, current geometry compression methods primarily focus on individual 3D models or sequences
27 of 3D models that are temporally correlated, but struggle to handle more general data sets, such as
28 compressing large numbers of unrelated 3D shapes.

29 Unlike images and videos represented as *regular* 2D or 3D volumes, mesh models are commonly
30 represented as triangle meshes, which are irregular and challenging to compress. Thus, a natural
31 idea is to structure the mesh models and then leverage image or video compression techniques to
32 compress them. Converting mesh models into voxelized point clouds is a common practice, and the
33 mesh models can be recovered from the point clouds via surface reconstruction methods [22, 24].
34 Based on this, in recent years, MPEG has developed two types of 3D point cloud compression (PCC)
35 standards [46, 28]: geometry-based PCC (GPCC) for static models and video-based PCC (VPCC) for
36 sequential models. And with advancements in deep learning, numerous learning-based PCC methods
37 [41, 14, 55, 19, 54] have emerged, enhancing compression efficiency. However, the voxelized point

38 clouds require a high resolution (typically 2^{10} or more) to accurately represent geometry data, which
39 is redundancy, limiting the compression efficiency.

40 Another regular representation involves utilizing implicit fields of mesh models, such as signed
41 distance fields (SDF) and truncated signed distance fields (TSDF). This is achieved by calculating
42 the value of the implicit field at each uniformly distributed grid point, resulting in a regular volume.
43 And the mesh models can be recovered from the implicit fields through Matching Cubes [32] or its
44 variants [15, 45]. Compared with point clouds, the implicit volume could represent the mesh models
45 in a relatively small resolution. Recently proposed methods, such as DeepSDF [36], utilize multilayer
46 perceptrons (MLPs) to regress the SDFs of any given query points. While this representation achieves
47 high accuracy for single or similar models (e.g., chairs, tables), the limited receptive field of MLPs
48 makes it challenging to represent large numbers of models in different categories, which is a more
49 common scenario in practice.

50 We propose NeCGS, a novel framework for compressing large sets of geometric models. Our NeCGS
51 framework consists of two stages: regular geometry representation and compact neural compression.
52 In the first stage, each model is converted into a regular 4D volumetric format, called the *TSDF-Def*
53 *volume*, which can be considered a 3D ‘image’. In the second stage, we use an auto-decoder to
54 regress these 4D volumes. The embedded features and decoder parameters represent these models,
55 and compressing these components allows us to compress the entire geometry set. We conducted
56 extensive experiments on various datasets, demonstrating that our NeCGS framework achieves higher
57 compression efficiency compared to existing geometry compression methods when handling large
58 numbers of models. Our NeCGS can achieve a compression ratio of nearly 900 on some datasets,
59 compressing hundreds or even thousands of different models into 1~2 MB while preserving detailed
60 structures.




Figure 1: Our NeCGeS can compress geometry data with hundreds or even thousands of shapes into 1~2 MB while preserving details. **Left:** Original Geometry Data. **Right:** Decompressed Geometry Data. [Q](#) Zoom in for details.

61 2 Related Work

62 2.1 Geometry Representation

63 In general, the representation of geometry data is divided into two main categories, explicit represen-
64 tation and implicit representation, and they could be transformed into another.

65 **Explicit Representation.** Among the explicit representations, voxelization [7] is the most intuitive.
66 In this method, geometry models are represented by regularly distributed grids, effectively converting
67 them into 3D ‘images’. While this approach simplifies the processing of geometry models using
68 image processing techniques, it requires a high resolution to accurately represent the models, which
69 demands substantial memory and limits its application. Another widely used geometry representation
70 method is the point cloud, which consists of discrete points sampled from the surfaces of models.
71 This method has become a predominant approach for surface representation [2, 39, 40]. However, the
72 discrete nature of the points imposes constraints on its use in downstream tasks such as rendering and
73 editing. Triangle meshes offer a more precise and efficient geometry representation. By approximating
74 surfaces with numerous triangles, they achieve higher accuracy and efficiency for certain downstream
75 tasks.

76 **Implicit Representation.** Implicit representations use the isosurface of a function or field to represent
77 surfaces. The most widely used implicit representations include Binary Occupancy Field (BOF)
78 [22, 35], Signed Distance Field (SDF) [36, 29], and Truncated Signed Distance Field (TSDF) [11],
79 from which the model’s surface can be easily extracted. However, these methods are limited to
80 representing watertight models. The Unsigned Distance Field (UDF) [8], which is the absolute value
81 of the SDF, can represent more general models, not just watertight ones. Despite this advantage,
82 extracting surfaces from UDF is challenging, which limits its application.

83 **Conversion between Geometry Representations.** Geometry representations can be converted
84 between explicit and implicit forms. Various methods [21, 22, 24, 6, 35, 29, 45] are available for
85 calculating the implicit field from given models. Conversely, when converting from implicit to
86 explicit forms, Marching Cubes [32] and its derivatives [48, 49, 15, 45] can reconstruct continuous
87 surfaces from various implicit fields.

88 2.2 3D Geometry Data Compression

89 **Single 3D Geometric Model Compression.** In recent decades, compression techniques for images
90 and videos have rapidly advanced [51, 34, 59, 5, 4]. However, the irregular nature of geometry
91 data makes it more challenging to compress compared to images and video, which are represented
92 as volumetric data. A natural approach is to convert geometry data into voxelized point clouds,
93 treating them as 3D ‘images’, and then applying image and video compression techniques to them.
94 Following this intuition, MPEG developed the GPCC standards [13, 28, 47], where triangle meshes or
95 triangle soup approximates the surfaces of 3D models, enabling the compression of models with more
96 complex structures. Subsequently, several improved methods [37, 60, 53, 62] and learning-based
97 methods [18, 43, 10, 9, 3, 42, 54] have been proposed to further enhance compression performance.
98 However, these methods rely on voxelized point clouds to represent geometry models, which is
99 inefficient and memory-intensive, limiting their compression efficiency. In contrast to the previously
100 mentioned methods, Draco [12] uses a kd-tree-based coding method to compress vertices and employs
101 the EdgeBreaker algorithm to encode the topological relationships of the geometry data. Draco
102 utilizes uniform quantization to control the compression ratio, but its performance decreases at higher
103 compression ratios.

104 **Multiple Model Compression.** Compared to compressing single 3D geometric models, compressing
105 multiple objects is significantly more challenging. SLRMA [17] addresses this by using a low-rank
106 matrix to approximate vertex matrices, thus compressing sequential models. Mekuria et al. [33]
107 proposed the first codec for compressing sequential point clouds, where each frame is coded using
108 Octree subdivision through an 8-bit occupancy code. Building on this concept, MPEG developed the
109 VPCC standards [13, 28, 47], which utilize 3D-to-2D projection and encode time-varying projected
110 planes, depth maps, and other data using video codecs. Several improved methods [57, 26, 1, 44]
111 have been proposed to enhance the compression of sequential models. Recently, shape priors like
112 SMPL [31] and SMAL [63] have been introduced, allowing the pose and shape of a template frame
113 to be altered using only a few parameters. Pose-driven geometry compression methods [16, 58, 56]

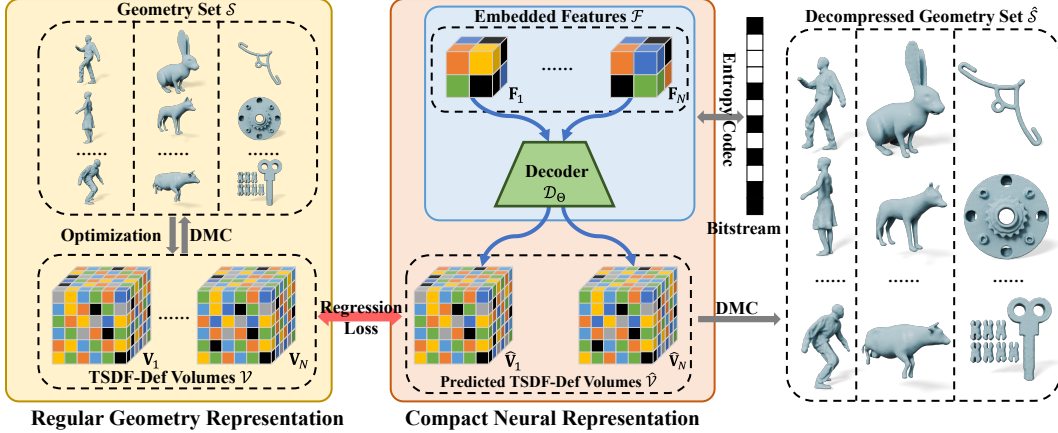


Figure 2: The pipeline of NeCGS. It first represents original meshes regularly into TSDF-Def volumes, and an auto-decoder network is utilized to regress these volume. Then the embedded features and decoder parameters are compressed into bitstreams through entropy coding. When decompressing the models, the decompressed embedded features are fed into the decoder with the decompressed parameters from the bitstreams, reconstructing the TSDF-Def volumes, and the models can be extracted from them.

114 leverage this approach to achieve high compression efficiency. However, these methods are limited to
 115 sequences of corresponding geometry data and cannot handle sets of unrelated geometry data, which
 116 is more common in practice.

117 3 Proposed Method

118 **Overview.** Given a set of N 3D mesh models containing diverse categories, denoted as $\mathcal{S} = \{\mathcal{S}_i\}_{i=1}^N$,
 119 we aim to compress them into a bitstream while maintaining the quality of the decompressed models
 120 as much as possible. To this end, we propose a neural compression paradigm called NeCGS. As
 121 shown in Fig. 2, NeCGS consists of two main modules, i.e., Regular Geometry Representation (RGR)
 122 and Compact Neural Representation (CNR). Specifically, RGR first represents each *irregular* mesh
 123 model within \mathcal{S} into a *regular* 4D volume, namely TSDF-Def volume that *mplicitly* describes the
 124 geometry structure of the model, via a rendering-based optimization, thus leading to a set of 4D
 125 volumes $\mathcal{V} := \{\mathbf{V}_i\}_{i=1}^N$ with \mathbf{V}_i corresponding to \mathcal{S}_i . Then CNR further obtains a more compact
 126 neural representation of \mathcal{V} , where a *quantization-aware* auto-decoder-based network is constructed
 127 to regress these volumes, producing an embedded feature for each volume. Finally, the embedded
 128 features along with the network parameters are encoded into a bitstream through a typical entropy
 129 coding method to achieve compression. We also want to **note** that NeCGS can also be applied to
 130 compress 3D geometry sets represented in *3D point clouds*, where one can either reconstruct from the
 131 given point clouds 3D surfaces through a typical surface reconstruction method or adopt a pre-trained
 132 network for SDF estimation from point clouds, e.g., SPSR [22] or IMLS [24], to bridge the gap
 133 between 3D mesh and point cloud models. In what follows, we will detail NeCGS.

134 3.1 Regular Geometry Representation

135 Unlike 2D images and videos, where pixels are uniformly
 136 distributed on 2D *regular* grids, the *irregular* characteristic
 137 of 3D mesh models makes it challenging to compress them
 138 efficiently and effectively. We propose to convert each
 139 3D mesh model to a 4D regular volume called TSDF-
 140 Def volume, which implicitly represents the geometry
 141 structure of the model. Such a regular representation can
 142 describe the model precisely, and its regular nature proves
 143 beneficial for compression in the subsequent stage.

144 **TSDF-Def Volume.** Although 3D regular SDF or TSDF
 145 volumes are widely used for representing 3D geometry
 146 models, they may introduce distortions when the volume




Figure 3: 2D visual illustration of DMC. The blue points refer to the deformable grid points, the green points refer to the vertices of the extracted surfaces, and the orange lines refer to the faces of the extracted surfaces. **Left:** The original grid points. **Right:** The surface extraction.

147 resolution is relatively limited. Inspired by recent shape extracting methods [48, 49], we propose
 148 TSDF-Def, which extends the regular TSDF volume by introducing an additional deformation for
 149 each grid point to adjust the detailed structure during the extraction of models, as shown in Fig.
 150 3. Accordingly, we develop the differentiable *Deformable Marching Cubes* (DMC), the variant of
 151 the Marching Cubes method [32], for surface extraction from a TSDF-Def volume. Consequently,
 152 each shape \mathbf{S} is represented as a 4D TSDF-Def volume, denoted as $\mathbf{V} \in \mathbb{R}^{K \times K \times K \times 4}$, where K
 153 is the volume resolution. More specifically, the value of the grid point located at (u, v, w) is
 154 $\mathbf{V}(u, v, w) := [\text{TSDF}(u, v, w), \Delta u, \Delta v, \Delta w]$, where $(\Delta u, \Delta v, \Delta w)$ are the deformation for the grid
 155 point and $1 \leq u, v, w \leq K$. TSDF-Def enhances representation accuracy, particularly when the grid
 156 resolution is relatively low.

157 **Optimization of TSDF-Def Volumes.** To obtain the optimal TSDF-Def volume \mathbf{V} for a given model
 158 \mathbf{S} , after initializing the deformations of each grid to zero and computing the TSDF value for each
 159 grid we optimize the following problem:

$$\min_{\mathbf{V}} \mathcal{E}_{\text{Rec}}(\text{DMC}(\mathbf{V}), \mathbf{S}), \quad (1)$$

160 where $\text{DMC}(\cdot)$ refers to the differentiable DMC process for extracting surfaces from TSDF-Def
 161 volumes, and the $\mathcal{E}_{\text{Reg}}(\cdot, \cdot)$ measures the differences between the rendered depth and silhouette
 162 images of two mesh models through the differentiable rasterization [25]. Algorithm 1 summarizes
 163 the whole optimization process. More details can be found in Sec. A.2 of the subsequent *Appendix*.

Algorithm 1: Optimization of TSDF-Def Volumes

Input: 3D mesh model \mathbf{S} ; the maximum number of iterations maxIter .

Output: The optimal TSDF-Def volume $\mathbf{V} \in \mathbb{R}^{K \times K \times K \times 4}$.

- 1 Place uniformly distributed grids in the cube of \mathbf{S} , denoted as $\mathbf{G} \in \mathbb{R}^{K \times K \times K \times 3}$;
 - 2 Initialize $\mathbf{V}[\dots, 0]$ as the ground truth TSDF of \mathbf{S} at the location of \mathbf{G} , the deformation
 $\mathbf{V}[\dots, 1:] = 0$, and the current iteration $\text{Iter} = 0$;
 - 3 **while** $\text{Iter} < \text{maxIter}$ **do**
 - 4 Recover shape from \mathbf{V} according to DMC, $\text{DMC}(\mathbf{V})$;
 - 5 Calculate the reconstruction error, $\mathcal{E}_{\text{Rec}}(\text{DMC}(\mathbf{V}), \mathbf{S})$;
 - 6 Optimize \mathbf{V} using ADAM optimizer based on the reconstruction error;
 - 7 $\text{Iter} := \text{Iter} + 1$;
 - 8 **end**
 - 9 **return** \mathbf{V} ;
-

164 3.2 Compact Neural Representation

165 Observing the similarity of local geometric structures within a typical 3D model and across different
 166 models, i.e., redundancy, we further propose a *quantization-aware* neural representation process
 167 to summarize the similarity within \mathcal{V} , leading to more compact representations with redundancy
 168 removed.

169 **Network Architecture.** We construct an auto-decoder network architecture to regress these 4D
 170 TSDF-Def volumes. Specifically, it is composed of a head layer, which increases the channel of its
 171 input, and L cascaded upsampling modules, which progressively upscale the feature volume. We
 172 also utilize the PixelShuffle technique [50] between the convolution and activation layers to achieve
 173 upscaling. We refer reviewers to Sec. B of *Appendix* for more details. For TSDF-Def volume \mathbf{V}_i ,
 174 the corresponding input to the auto-decoder is the embedded feature, denoted as $\mathbf{F}_i \in \mathbb{R}^{K' \times K' \times K' \times C}$,
 175 where K' is the resolution satisfying $K' \ll K$ and C is the number of channels. Moreover, we
 176 integrate differentiable quantization to the embedded features and network parameters in the process,
 177 which can efficiently reduce the quantization error. In all, the compact neural representation process
 178 can be written as

$$\widehat{\mathbf{V}}_i = \mathcal{D}_{\mathcal{Q}(\Theta)}(\mathcal{Q}(\mathbf{F}_i)). \quad (2)$$

179 where $\mathcal{Q}(\cdot)$ stands for the differentiable quantization operator, and $\widehat{\mathbf{V}}_i$ is the regressed TSDF-Def.

180 **Loss Function.** We employ a joint loss function comprising Mean Absolute Error (MAE) and
 181 Structural Similarity Index (SSIM) to simultaneously optimize the embedded features $\{\mathbf{F}_i\}$ and

182 the network parameters Θ . In computing the MAE between the predicted and ground truth TSDF-
 183 Def volumes, we concentrate more on the grids close to the surface. These surface grids crucially
 184 determine the surfaces through their TSDFs and deformations; hence we assign them higher weights
 185 during optimization than the grids farther away from the surface. The overall loss function for the
 186 i -th model is written as

$$\mathcal{L}(\widehat{\mathbf{V}}_i, \mathbf{V}_i) = \|\widehat{\mathbf{V}}_i - \mathbf{V}_i\|_1 + \lambda_1 \|\mathbf{M}_i \odot (\widehat{\mathbf{V}}_i - \mathbf{V}_i)\|_1 + \lambda_2 (1 - \text{SSIM}(\widehat{\mathbf{V}}_i, \mathbf{V}_i)), \quad (3)$$

187 where $\mathbf{M}_i = \mathbb{1}(|\mathbf{V}_i[\dots, 0]| < \tau)$ is the mask, indicating whether a grid is near the surface, i.e., its
 188 TSDF is less than the threshold τ , while λ_1 and λ_2 are the weights to balance each term of the loss
 189 function.

190 **Entropy Coding.** After obtaining the quantized features $\{\tilde{\mathbf{F}}_i = \mathcal{Q}(\mathbf{F}_i)\}$ and quantized network
 191 parameters $\tilde{\Theta} = \mathcal{Q}(\Theta)$, we adopt the Huffman Codec [20] to further compress them into a bit-
 192 stream. More advanced entropy coding methods can be employed to further improve compression
 193 performance.

194 3.3 Decompression

195 To obtain the 3D mesh models from the bitstream, we first decompress the bitstream to derive the
 196 embedded features, $\{\tilde{\mathbf{F}}_i\}$ and the decoder parameter, $\tilde{\Theta}$. Then, for each $\tilde{\mathbf{F}}_i$, we feed it to the decoder
 197 $\mathcal{D}_{\tilde{\Theta}}(\cdot)$ to generate its corresponding TSDF-Def volume

$$\widehat{\mathbf{V}}_i = \mathcal{D}_{\tilde{\Theta}}(\tilde{\mathbf{F}}_i). \quad (4)$$

198 Finally, we utilize DMC to recover each shape from $\widehat{\mathbf{V}}_i$, $\widehat{\mathbf{S}}_i = \text{DMC}(\widehat{\mathbf{V}}_i)$, forming the set of decom-
 199 pressed geometry data, $\widehat{\mathcal{S}} = \{\widehat{\mathbf{S}}_i\}_{i=1}^N$.

200 4 Experiment

201 4.1 Experimental Setting

202 **Implementation details.** In the process of optimizing TSDF-Def volumes, we employed the ADAM
 203 optimizer [23] for 500 iterations per shape, using a learning rate of 0.01. The resolution of TSDF-Def
 204 volumes was $K = 128$. The resolution and the number of channels of the embedded features were
 205 $K' = 4$ and $C = 16$, respectively. And the decoder is composed of $L = 5$ upsampling modules with
 206 an up-scaling factor of 2. During the optimization, we set $\lambda_1 = 5$ and $\lambda_2 = 10$, and the embedded
 207 features and decoder parameters were optimized by the ADAM optimizer for 400 epochs, with a
 208 learning rate of 1e-3. We achieved different compression efficiencies by adjusting decoder sizes. We
 209 conducted all experiments on an NVIDIA RTX 3090 GPU with Intel(R) Xeon(R) CPU.

210 **Datasets.** We tested our NeCGS on various types
 211 of datasets, including humans, animals, and CAD
 212 models. For human models, we randomly selected
 213 500 shapes from the AMA dataset [52]. For animal
 214 models, we randomly selected 500 shapes from
 215 the DT4D dataset [27]. For the CAD models, we
 216 randomly selected 1000 shapes from the Thingi10K
 217 dataset [61]. Besides, we randomly selected 200
 218 models from each dataset, forming a more challenging dataset, denoted as Mixed. The details about
 219 the selected datasets are shown in Table 1. In all experiments, we scaled all models in a cube with a
 220 range of $[-1, 1]^3$ to ensure they are in the same scale.

222 **Methods under Comparison.** In terms of traditional geometry codecs, we chose the three most
 223 impactful geometry coding standards with released codes, G-PCC² and V-PCC³ from MPEG (see

Table 1: Details of the selected datasets¹.

Dataset	Original Size (MB)	# Models
AMA	378.41	500
DT4D	683.80	500
Thingi10K	335.92	1000
Mixed	496.16	600

¹The original geometry data is kept as triangle meshes, so the storage size is much less than the voxelized point clouds.

²<https://github.com/MPEGGroup/mpeg-pcc-tmc13>

³<https://github.com/MPEGGroup/mpeg-pcc-tmc2>

224 more details about them in [13, 28, 47]), and Draco⁴ from Google as the baseline methods. Addi-
 225 tionally, we compared our approach with state-of-the-art deep learning-based compression methods,
 226 specifically PCGCv2 [54]. Furthermore, we adapted DeepSDF [36] with quantization to serve as
 227 another baseline method, denoted as QuantDeepSDF. It is worth noting that while some of the chosen
 228 baseline methods were originally designed for point cloud compression, we utilized voxel sampling
 229 and SPSR [22] to convert them between the forms of point cloud and surface. More details can be
 230 found in Sec. C.2 appendix.

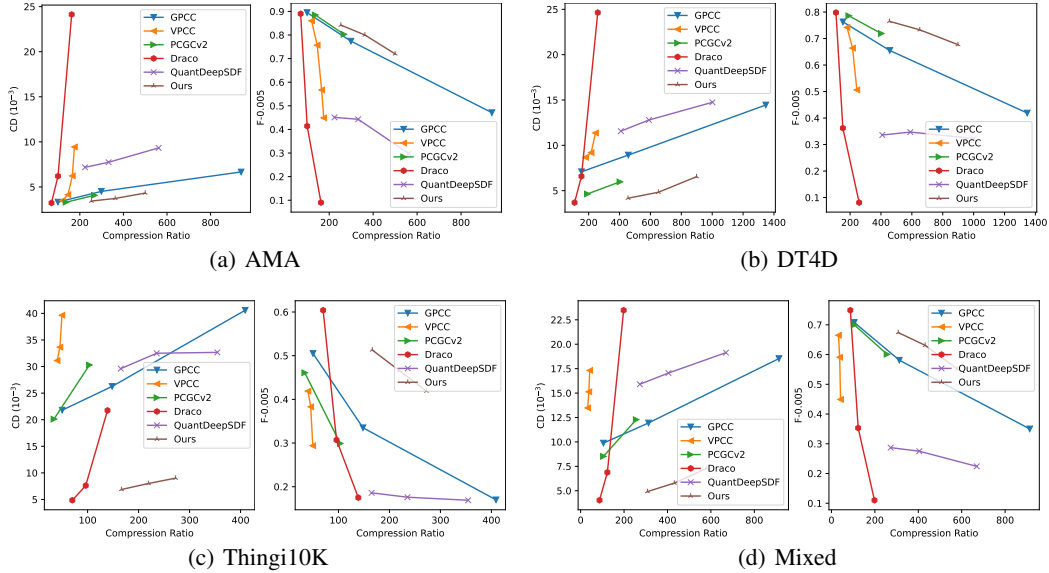


Figure 4: Quantitative comparisons of different methods on four 3D geometry sets.

231 **Evaluation Metrics.** Following previous reconstruction methods [35, 38], we utilize Chamfer
 232 Distance (CD), Normal Consistency (NC), F-Score with the thresholds of 0.005 and 0.01 (F1-0.005
 233 and F1-0.01) as the evaluation metrics. Furthermore, to comprehensively compare the compression
 234 efficiency of different methods, we use Rate-Distortion (RD) curves. These curves illustrate the
 235 distortions at various compression ratios, with CD and F1-0.005 specifically describing the distortion
 236 of the decompressed models. Our goal is to minimize distortion, indicated by a low CD and a high
 237 F1-Score, while maximizing the compression ratio. Therefore, for the RD curve representing CD,
 238 optimal compression performance is achieved when the curve is closest to the lower right corner.
 239 Similarly, for the RD curve representing the F1-Score, the ideal compression performance is when
 240 the curve is nearest to the upper right corner. Their detailed definition can be found in Sec. C.1 of
 241 appendix.

242 4.2 Results

243 The RD curves of different compression methods under different datasets are shown in Fig. 4. As
 244 the compression ratio increases, the distortion also becomes larger. It is obvious that our NeCGS
 245 can achieve much better compression performance than the baseline methods when the compression
 246 ratio is high, even in the challenging Mixed dataset. In particular, our NeCGS achieves a minimum
 247 compression ratio of 300, and on the DT4D dataset, the compression ratio even reaches nearly 900,
 248 with minimal distortion. Due to the larger model differences within the Thingi10K and Mixed
 249 datasets compared to the other two datasets, the compression performance on these two datasets is inferior.

250 The visual results of different compression meth-
 251 ods are shown in Fig. 5. Compared to other
 252 methods, models compressed using our ap-
 253 proach occupy a larger compression ratio and
 254 retain more details after decompression. Fig. 6


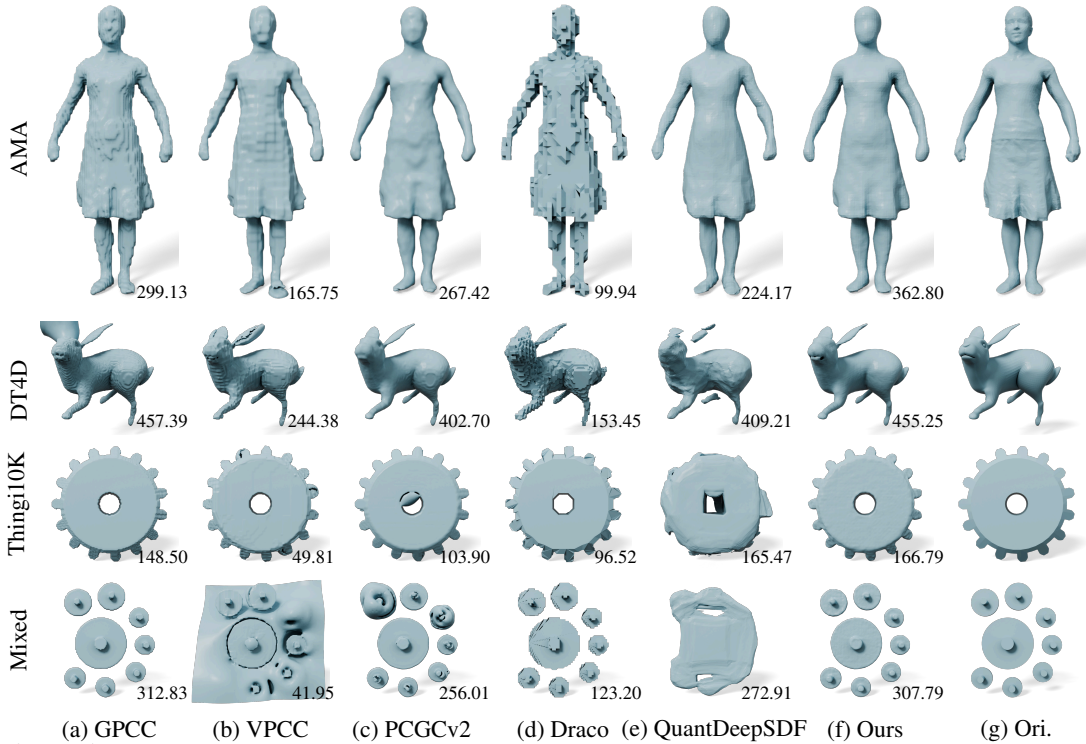


Figure 6: Decompressed models under different compression ratios.

⁴<https://github.com/google/draco>



(a) GPCC (b) VPCC (c) PCGCv2 (d) Draco (e) QuantDeepSDF (f) Ours (g) Ori.
 Figure 5: Visual comparisons of different compression methods. All numbers in corners represent the compression ratio. [Zoom in for details.](#)

255 illustrates the decompressed models under different compression ratio. Even when the compression
 256 ratio reaches nearly 900, our method can still retain the details of the models.

257 4.3 Ablation Study

258 In order to illustrate the efficiency of each design of our NeCGS, we conducted extensive ablation
 259 study about them on the Mixed dataset.

260 Necessity of the Deformation of

261 **Grids.** We utilize TSDF-Def volumes
 262 to as the regular geometry representa-
 263 tion, instead of TSDF volumes like
 264 previous methods. Compared with
 265 models recovered from TSDF vol-
 266 umes through MC, the models recovered
 267 from TSDF-Def volumes through
 268 DMC preserve more details of the thin
 269 structures, especially when the volume resolutions are relatively small, as shown in Fig. 7. We also
 270 conducted a numerical comparison of the decompressed models on the AMA dataset under these two
 271 settings, and the results are shown in Table. 2, demonstrating its advantages.



Figure 7: Models recovered from different regular geometry representations under various volume resolutions. From **Left to Right**: Original, TSDF with $K = 64$, TSDF with $K = 128$, TSDF-Def with $K = 64$, and TSDF-Def with $K = 128$.

Table 2: Quantitative comparisons of different RGRs.

RGR	Size (MB)	Com. Ratio	CD ($\times 10^{-3}$) \downarrow	NC \uparrow	F1-0.005 \uparrow	F1-0.01 \uparrow
TSDF	1.631	304.20	5.015	0.944	0.662	0.936
TSDF-Def	1.612	307.79	4.913	0.947	0.674	0.943

272 **Neural Representation Structure.** To illustrate the superiority of auto-decoder framework, we
 273 utilize an auto-encoder to regress the TSDF-Def volume. Technically, we used a ConvNeXt block
 274 [30] as the encoder by replacing 2D convolutions with 3D convolutions. Under the auto-encoder
 275 framework, we optimize the parameters of the encoder to change the embedded features. The RD




Figure 8: (a) RD curves of different neural representation structures. (b) RD curves of different regression losses.

276 curves about these two structures are shown in Fig. 8(a), demonstrating rationality of our decoder
 277 structure.

278 **SSIM Loss.** Compared to MAE, which focuses on
 279 one-to-one errors between predicted and ground truth
 280 volumes, the SSIM item in Eq. 3 emphasizes more
 281 on the local similarity between volumes, increasing
 282 the regression accuracy. To verify this, we removed
 283 the SSIM item and kept others unchanged. Their RD
 284 curves are shown in Fig. 8(b), and it is obvious that
 285 the SSIM item in the regression loss increases the
 286 compression performance. The visual comparison is
 287 shown in Fig. 9, and without SSIM, there are floating
 288 parts around the decompressed models.




Figure 9: Visual comparison of regression loss w/ and w/o SSIM item.

289 **Resolution of TSDF-Def Volumes.** We tested the com-
 290 pression performance at different resolutions of TSDF-
 291 Def volumes by adjusting the decoder layers accordingly.
 292 Specifically, we removed the last layer for a resolution
 293 of 64 and added an extra layer for a resolution of 256.
 294 The quantitative and numerical comparisons are shown in
 295 Table 3 and Fig. 10, respectively. Obviously, increasing
 296 the volume resolution can enhance the compression effec-
 297 tiveness, resulting in more detailed structures preserved
 298 after decompression. However, the optimization and infer-
 299 ence time also increase accordingly due to more layers
 300 involved.




Figure 10: Visual comparison under different resolutions of TSDF-Def volume.

Table 3: Quantitative comparisons of different resolutions of TSDF-Def volumes.

Res.	Size (MB)	Com. Ratio	CD ($\times 10^{-3}$) ↓	NC ↑	F1-0.005 ↑	F1-0.01 ↑	Opt Time (h)	Infer. Time (ms)
64	1.408	268.75	4.271	0.927	0.721	0.966	2.16	38.97
128	1.493	253.45	3.436	0.952	0.842	0.991	16.32	98.95
256	1.627	232.58	3.234	0.962	0.870	0.995	94.50	421.94

301 5 Conclusion and Discussion

302 We have presented NeCGS, a highly effective neural compression scheme for 3D geometry sets.
 303 NeCGS has achieved remarkable compression performance on various datasets with diverse and
 304 detailed shapes, outperforming state-of-the-art compression methods to a large extent. These advan-
 305 tages are attributed to our regular geometry representation and the compression accomplished by a
 306 convolution-based auto-decoder. We believe our NeCGS framework will inspire further advancements
 307 in the field of geometry compression.

308 However, our method still suffers from the following two limitations. One is that it requires more
 309 than 15 hours to regress the TSDF-Def volumes, and the other one is that the usage of 3D convolution
 310 layers limits the inference speed. Our future work will focus on addressing these challenges by
 311 accelerating the optimization process and incorporating more efficient network modules.

References

- 312
- 313 [1] A. Ahmmed, M. Paul, M. Murshed, and D. Taubman. Dynamic point cloud geometry compression using
314 cuboid based commonality modeling framework. In *2021 IEEE International Conference on Image
315 Processing (ICIP)*, pages 2159–2163. IEEE, 2021. 3
- 316 [2] P. J. Besl and N. D. McKay. Method for registration of 3-d shapes. In *Sensor Fusion IV: Control Paradigms
317 and Data Structures*, volume 1611, pages 586–606. Spie, 1992. 3
- 318 [3] S. Biswas, J. Liu, K. Wong, S. Wang, and R. Urtasun. Muscle: Multi sweep compression of lidar using
319 deep entropy models. *Advances in Neural Information Processing Systems*, 33:22170–22181, 2020. 3
- 320 [4] H. Chen, M. Gwilliam, S.-N. Lim, and A. Shrivastava. Hnerv: A hybrid neural representation for
321 videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages
322 10270–10279, 2023. 3
- 323 [5] H. Chen, B. He, H. Wang, Y. Ren, S. N. Lim, and A. Shrivastava. Nerv: Neural representations for videos.
324 *Advances in Neural Information Processing Systems*, 34:21557–21568, 2021. 3
- 325 [6] Z.-Q. Cheng, Y.-Z. Wang, B. Li, K. Xu, G. Dang, and S.-Y. Jin. A survey of methods for moving least
326 squares surfaces. In *Proceedings of the Fifth Eurographics/IEEE VGTC conference on Point-Based
327 Graphics*, pages 9–23, 2008. 3
- 328 [7] J. Chibane, T. Alldieck, and G. Pons-Moll. Implicit functions in feature space for 3d shape reconstruction
329 and completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*,
330 pages 6970–6981, June 2020. 3
- 331 [8] J. Chibane, G. Pons-Moll, et al. Neural unsigned distance fields for implicit function learning. *Advances in
332 Neural Information Processing Systems*, 33:21638–21652, 2020. 3
- 333 [9] T. Fan, L. Gao, Y. Xu, D. Wang, and Z. Li. Multiscale latent-guided entropy model for lidar point cloud
334 compression. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(12):7857–7869, 2023.
335 3
- 336 [10] C. Fu, G. Li, R. Song, W. Gao, and S. Liu. Octattention: Octree-based large-scale contexts model for point
337 cloud compression. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages
338 625–633, 2022. 3
- 339 [11] P. Gao, Z. Jiang, H. You, P. Lu, S. C. Hoi, X. Wang, and H. Li. Dynamic fusion with intra-and inter-modality
340 attention flow for visual question answering. In *Proceedings of the IEEE/CVF conference on computer
341 vision and pattern recognition*, pages 6639–6648, 2019. 3
- 342 [12] Google. Point cloud compression reference software. Website. <https://github.com/google/draco>. 3
- 343 [13] D. Graziosi, O. Nakagami, S. Kuma, A. Zaghetto, T. Suzuki, and A. Tabatabai. An overview of ongoing
344 point cloud compression standardization activities: Video-based (v-pcc) and geometry-based (g-pcc).
345 *APSIPA Transactions on Signal and Information Processing*, 9:e13, 2020. 3, 7
- 346 [14] A. F. Guarda, N. M. Rodrigues, and F. Pereira. Point cloud coding: Adopting a deep learning-based
347 approach. In *2019 Picture Coding Symposium (PCS)*, pages 1–5. IEEE, 2019. 1
- 348 [15] B. Guillard, F. Stella, and P. Fua. Meshudf: Fast and differentiable meshing of unsigned distance field
349 networks. In *European Conference on Computer Vision*, pages 576–592, 2022. 2, 3
- 350 [16] J. Hou, L.-P. Chau, N. Magnenat-Thalmann, and Y. He. Compressing 3-d human motions via keyframe-
351 based geometry videos. *IEEE Transactions on Circuits and Systems for Video Technology*, 25(1):51–62,
352 2014. 3
- 353 [17] J. Hou, L.-P. Chau, N. Magnenat-Thalmann, and Y. He. Sparse low-rank matrix approximation for data
354 compression. *IEEE Transactions on Circuits and Systems for Video Technology*, 27(5):1043–1054, 2015. 3
- 355 [18] L. Huang, S. Wang, K. Wong, J. Liu, and R. Urtasun. Otsqueeze: Octree-structured entropy model for
356 lidar compression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,
357 pages 1313–1323, 2020. 3
- 358 [19] T. Huang and Y. Liu. 3d point cloud geometry compression on deep learning. In *Proceedings of the 27th
359 ACM international conference on multimedia*, pages 890–898, 2019. 1
- 360 [20] D. A. Huffman. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*,
361 40(9):1098–1101, 1952. 6

- 362 [21] M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth*
363 *Eurographics symposium on Geometry processing*, pages 61–70, 2006. 3
- 364 [22] M. Kazhdan and H. Hoppe. Screened poisson surface reconstruction. *ACM Transactions on Graphics*
365 *(ToG)*, 32(3):1–13, 2013. 1, 3, 4, 7
- 366 [23] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*,
367 2014. 6
- 368 [24] R. Kolluri. Provably good moving least squares. *ACM Transactions on Algorithms*, 4(2):1–25, 2008. 1, 3,
369 4
- 370 [25] S. Laine, J. Hellsten, T. Karras, Y. Seol, J. Lehtinen, and T. Aila. Modular primitives for high-performance
371 differentiable rendering. *ACM Transactions on Graphics (ToG)*, 39(6):1–14, 2020. 5
- 372 [26] L. Li, Z. Li, V. Zakharchenko, J. Chen, and H. Li. Advanced 3d motion prediction for video-based dynamic
373 point cloud compression. *IEEE Transactions on Image Processing*, 29:289–302, 2019. 3
- 374 [27] Y. Li, H. Takehara, T. Taketomi, B. Zheng, and M. Nießner. 4dcomplete: Non-rigid motion estimation
375 beyond the observable surface. In *Proceedings of the IEEE/CVF International Conference on Computer*
376 *Vision*, pages 12706–12716, 2021. 6
- 377 [28] H. Liu, H. Yuan, Q. Liu, J. Hou, and J. Liu. A comprehensive study and comparison of core technologies
378 for mpeg 3-d point cloud compression. *IEEE Transactions on Broadcasting*, 66(3):701–717, 2019. 1, 3, 7
- 379 [29] S.-L. Liu, H.-X. Guo, H. Pan, P.-S. Wang, X. Tong, and Y. Liu. Deep implicit moving least-squares
380 functions for 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
381 *Pattern Recognition*, pages 1788–1797, June 2021. 3
- 382 [30] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie. A convnet for the 2020s. In *Proceedings*
383 *of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11976–11986, 2022. 8
- 384 [31] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. Smpl: A skinned multi-person linear
385 model. *ACM Trans. Graph.*, 34(6), oct 2015. 3
- 386 [32] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3d surface construction algorithm.
387 *ACM siggraph computer graphics*, 21(4):163–169, 1987. 2, 3, 5
- 388 [33] R. Mekuria, K. Blom, and P. Cesar. Design, implementation, and evaluation of a point cloud codec for
389 tele-immersive video. *IEEE Transactions on Circuits and Systems for Video Technology*, 27(4):828–842,
390 2016. 3
- 391 [34] F. Mentzer, E. Agustsson, M. Tschanen, R. Timofte, and L. V. Gool. Practical full resolution learned
392 lossless image compression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*
393 *recognition*, pages 10629–10638, 2019. 3
- 394 [35] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger. Occupancy networks: Learning 3d
395 reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
396 *Pattern Recognition*, pages 4460–4470, June 2019. 3, 7
- 397 [36] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove. Deepsdf: Learning continuous signed
398 distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer*
399 *Vision and Pattern Recognition*, pages 165–174, June 2019. 2, 3, 7
- 400 [37] E. Peixoto. Intra-frame compression of point cloud geometry using dyadic decomposition. *IEEE Signal*
401 *Processing Letters*, 27:246–250, 2020. 3
- 402 [38] S. Peng, M. Niemeyer, L. Mescheder, M. Pollefeys, and A. Geiger. Convolutional occupancy networks. In
403 *European Conference on Computer Vision*, pages 523–540. Springer, 2020. 7
- 404 [39] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. Pointnet: Deep learning on point sets for 3d classification and
405 segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages
406 652–660, 2017. 3
- 407 [40] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a
408 metric space. *Advances in neural information processing systems*, 30:1–xxx, 2017. 3
- 409 [41] M. Quach, G. Valenzise, and F. Dufaux. Learning convolutional transforms for lossy point cloud geometry
410 compression. In *2019 IEEE international conference on image processing (ICIP)*, pages 4320–4324. IEEE,
411 2019. 1

- 412 [42] M. Quach, G. Valenzise, and F. Dufaux. Learning convolutional transforms for lossy point cloud geometry
413 compression. In *2019 IEEE international conference on image processing (ICIP)*, pages 4320–4324. IEEE,
414 2019. 3
- 415 [43] Z. Que, G. Lu, and D. Xu. Voxelcontext-net: An octree based framework for point cloud compression. In
416 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6042–6051,
417 2021. 3
- 418 [44] E. Ramalho, E. Peixoto, and E. Medeiros. Silhouette 4d with context selection: Lossless geometry
419 compression of dynamic point clouds. *IEEE Signal Processing Letters*, 28:1660–1664, 2021. 3
- 420 [45] S. Ren, J. Hou, X. Chen, Y. He, and W. Wang. Geoudf: Surface reconstruction from 3d point clouds via
421 geometry-guided distance representation. In *Proceedings of the IEEE/CVF International Conference on
422 Computer Vision*, pages 14214–14224, 2023. 2, 3
- 423 [46] S. Schwarz, M. Preda, V. Baroncini, M. Budagavi, P. Cesar, P. A. Chou, R. A. Cohen, M. Krivokuća,
424 S. Lasserre, Z. Li, et al. Emerging mpeg standards for point cloud compression. *IEEE Journal on Emerging
425 and Selected Topics in Circuits and Systems*, 9(1):133–148, 2018. 1
- 426 [47] S. Schwarz, M. Preda, V. Baroncini, M. Budagavi, P. Cesar, P. A. Chou, R. A. Cohen, M. Krivokuća,
427 S. Lasserre, Z. Li, et al. Emerging mpeg standards for point cloud compression. *IEEE Journal on Emerging
428 and Selected Topics in Circuits and Systems*, 9(1):133–148, 2018. 3, 7
- 429 [48] T. Shen, J. Gao, K. Yin, M.-Y. Liu, and S. Fidler. Deep marching tetrahedra: a hybrid representation for
430 high-resolution 3d shape synthesis. *Advances in Neural Information Processing Systems*, 34:6087–6101,
431 2021. 3, 5
- 432 [49] T. Shen, J. Munkberg, J. Hasselgren, K. Yin, Z. Wang, W. Chen, Z. Gojcic, S. Fidler, N. Sharp, and J. Gao.
433 Flexible isosurface extraction for gradient-based mesh optimization. *ACM Transactions on Graphics
434 (TOG)*, 42(4):1–16, 2023. 3, 5
- 435 [50] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang. Real-time
436 single image and video super-resolution using an efficient sub-pixel convolutional neural network. In
437 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1874–1883, 2016.
438 5
- 439 [51] Y. Strümpfer, J. Postels, R. Yang, L. V. Gool, and F. Tombari. Implicit neural representations for image
440 compression. In *European Conference on Computer Vision*, pages 74–91. Springer, 2022. 3
- 441 [52] D. Vlastic, I. Baran, W. Matusik, and J. Popović. Articulated mesh animation from multi-view silhouettes.
442 *ACM Transactions on Graphics*, 27(3):1–9, 2008. 6
- 443 [53] C. Wang, W. Zhu, Y. Xu, Y. Xu, and L. Yang. Point-voting based point cloud geometry compression. In
444 *2021 IEEE 23rd International Workshop on Multimedia Signal Processing (MMSP)*, pages 1–5. IEEE,
445 2021. 3
- 446 [54] J. Wang, D. Ding, Z. Li, and Z. Ma. Multiscale point cloud geometry compression. In *2021 Data
447 Compression Conference (DCC)*, pages 73–82. IEEE, 2021. 1, 3, 7
- 448 [55] J. Wang, H. Zhu, H. Liu, and Z. Ma. Lossy point cloud geometry compression via end-to-end learning.
449 *IEEE Transactions on Circuits and Systems for Video Technology*, 31(12):4909–4923, 2021. 1
- 450 [56] X. Wu, P. Zhang, M. Wang, P. Chen, S. Wang, and S. Kwong. Geometric prior based deep human point
451 cloud geometry compression. *IEEE Transactions on Circuits and Systems for Video Technology*, 2024. 3
- 452 [57] J. Xiong, H. Gao, M. Wang, H. Li, K. N. Ngan, and W. Lin. Efficient geometry surface coding in v-pcc.
453 *IEEE Transactions on Multimedia*, 25:3329–3342, 2022. 3
- 454 [58] R. Yan, Q. Yin, X. Zhang, Q. Zhang, G. Zhang, and S. Ma. Pose-driven compression for dynamic 3d
455 human via human prior models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 3
- 456 [59] Y. Yang, R. Bamler, and S. Mandt. Improving inference for neural image compression. *Advances in Neural
457 Information Processing Systems*, 33:573–584, 2020. 3
- 458 [60] X. Zhang, W. Gao, and S. Liu. Implicit geometry partition for point cloud compression. In *2020 Data
459 Compression Conference (DCC)*, pages 73–82. IEEE, 2020. 3
- 460 [61] Q. Zhou and A. Jacobson. Thingi10k: A dataset of 10,000 3d-printing models. *arXiv preprint
461 arXiv:1605.04797*, 2016. 6

- 462 [62] W. Zhu, Y. Xu, D. Ding, Z. Ma, and M. Nilsson. Lossy point cloud geometry compression via region-wise
463 processing. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(12):4575–4589, 2021. 3
- 464 [63] S. Zuffi, A. Kanazawa, D. Jacobs, and M. J. Black. 3D menagerie: Modeling the 3D shape and pose of
465 animals. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 3

466 Appendix

467 A Regular Geometry Representation

468 A.1 Tensor Quantization

469 Denoted \mathbf{x} is a tensor, we quantize it in a fixed interval, $[a, b]$, at $(2^N + 1)$ levels⁵ by

$$\mathcal{Q}(\mathbf{x}) = \text{Round} \left(\frac{\text{Clamp}(\mathbf{x}, a, b) - a}{s} \right) \times s + a, \quad (5)$$

470 where $s = (b - a)/2^N$. In our experiment, we set $a = -1$ and $b = 1$.

471 A.2 Optimization of TSDF-deformation Volumes

472 We set a series of camera pose, $\mathcal{T} = \{\mathbf{T}_i\}_{i=1}^E$, around the meshes. Let $\mathbf{I}_1^P(\mathbf{T}_i)$ and $\mathbf{I}_2^P(\mathbf{T}_i)$ represent
473 the depth images obtained from the reconstructed mesh $\text{DMC}(\mathbf{V})$ and the given mesh \mathbf{S} at the pose \mathbf{T}_i
474 respectively. Similarly, let $\mathbf{I}_1^M(\mathbf{T}_i)$ and $\mathbf{I}_2^M(\mathbf{T}_i)$ denote their respective silhouette images at pose \mathbf{T}_i .
475 The reconstruction error produced by silhouette and depth images at all pose are

$$\mathcal{E}_M(\text{DMC}(\mathbf{V}), \mathbf{S}) = \sum_{\mathbf{T}_i \in \mathcal{T}} \|\mathbf{I}_1^M(\mathbf{T}_i) - \mathbf{I}_2^M(\mathbf{T}_i)\|_1 \quad (6)$$

476 and

$$\mathcal{E}_D(\text{DMC}(\mathbf{V}), \mathbf{S}) = \sum_{\mathbf{T}_i \in \mathcal{T}} \|\mathbf{I}_1^P(\mathbf{T}_i) - \mathbf{I}_2^P(\mathbf{T}_i)\|_1 * \mathbf{I}_2^M(\mathbf{T}_i) \quad (7)$$

477 Then the reconstruction error is defined as

$$\mathcal{E}_{\text{Rec}}(\text{DMC}(\mathbf{V}), \mathbf{S}) = \mathcal{E}_M(\text{DMC}(\mathbf{V}), \mathbf{S}) + \lambda_{\text{rec}} \mathcal{E}_D(\text{DMC}(\mathbf{V}), \mathbf{S}), \quad (8)$$

478 where $E = 4$ and $\lambda_{\text{rec}} = 10$ in our experiment.

479 B Auto-decoder-based Neural Compression

480 B.1 Upsampling Module

481 In each upsampling module, we utilize a PixelShuffle layer between the convolution and activa-
482 tion layers to upscale the input, as shown in Fig. 11. The input feature volume has dimensions
483 $(N_{\text{in}}, N_{\text{in}}, N_{\text{in}}, C_{\text{in}})$, with an upsampling scale of s and an output channel count of C_{out} .

484 C Experiment

485 C.1 Evaluation Metric

486 Let \mathbf{S}_{Rec} and \mathbf{S}_{GT} denote the reconstructed and ground-truth 3D shapes, respectively. We then
487 randomly sample $N_{\text{eval}} = 10^5$ points on them, obtaining two point clouds, \mathbf{P}_{Rec} and \mathbf{P}_{GT} . For each
488 point of \mathbf{P}_{Rec} and \mathbf{P}_{GT} , the normal of the triangle face where it is sampled is considered to be its
489 normal vector, and the normal sets of \mathbf{P}_{Rec} and \mathbf{P}_{GT} are denoted as \mathbf{N}_{Rec} and \mathbf{N}_{GT} , respectively.
490 Let $\text{NN_Point}(\mathbf{x}, \mathbf{P})$ be the operator that returns the nearest point of \mathbf{x} in the point cloud \mathbf{P} . The CD
491 between them is defined as

$$\begin{aligned} \text{CD}(\mathbf{S}_{\text{Rec}}, \mathbf{S}_{\text{GT}}) &= \frac{1}{2N_{\text{eval}}} \sum_{\mathbf{x} \in \mathbf{P}_{\text{Rec}}} \|\mathbf{x} - \text{NN_Point}(\mathbf{x}, \mathbf{P}_{\text{GT}})\|_2 \\ &\quad + \frac{1}{2N_{\text{eval}}} \sum_{\mathbf{x} \in \mathbf{P}_{\text{GT}}} \|\mathbf{x} - \text{NN_Point}(\mathbf{x}, \mathbf{P}_{\text{Rec}})\|_2. \end{aligned} \quad (9)$$

⁵We partition the interval $[a, b]$ into $(2^N + 1)$ levels, rather than 2^N levels, to ensure the inclusion of the value 0.

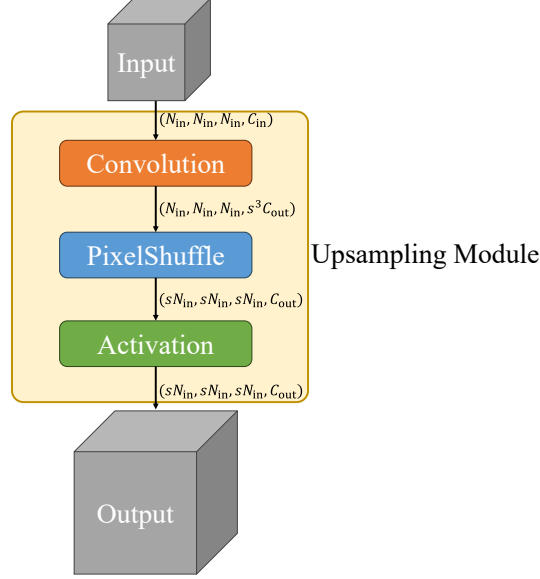


Figure 11: Upsampling Module.

492 Let $\text{NN_Normal}(\mathbf{x}, \mathbf{P})$ be the operator that returns the normal vector of the point \mathbf{x} 's nearest point in
 493 the point cloud \mathbf{P} . The NC is defined as

$$\begin{aligned} \text{NC}(\mathbf{S}_{\text{Rec}}, \mathbf{S}_{\text{GT}}) &= \frac{1}{2N_{\text{eval}}} \sum_{\mathbf{x} \in \mathbf{P}_{\text{Rec}}} |\mathbf{N}_{\text{Rec}}(\mathbf{x}) \cdot \text{NN_Normal}(\mathbf{x}, \mathbf{P}_{\text{GT}})| \\ &+ \frac{1}{2N_{\text{eval}}} \sum_{\mathbf{x} \in \mathbf{P}_{\text{GT}}} |\mathbf{N}_{\text{GT}}(\mathbf{x}) \cdot \text{NN_Normal}(\mathbf{x}, \mathbf{P}_{\text{Rec}})|. \end{aligned} \quad (10)$$

494 F-Score is defined as the harmonic mean between the precision and the recall of points that lie within
 495 a certain distance threshold ϵ between \mathbf{S}_{Rec} and \mathbf{S}_{GT} ,

$$\text{F-Score}(\mathbf{S}_{\text{Rec}}, \mathbf{S}_{\text{GT}}, \epsilon) = \frac{2 \cdot \text{Recall} \cdot \text{Precision}}{\text{Recall} + \text{Precision}}, \quad (11)$$

496 where

$$\begin{aligned} \text{Recall}(\mathbf{S}_{\text{Rec}}, \mathbf{S}_{\text{GT}}, \epsilon) &= \left| \left\{ \mathbf{x}_1 \in \mathbf{P}_{\text{Rec}}, \text{s.t. } \min_{\mathbf{x}_2 \in \mathbf{P}_{\text{GT}}} \|\mathbf{x}_1 - \mathbf{x}_2\|_2 < \epsilon \right\} \right|, \\ \text{Precision}(\mathbf{S}_{\text{Rec}}, \mathbf{S}_{\text{GT}}, \epsilon) &= \left| \left\{ \mathbf{x}_2 \in \mathbf{P}_{\text{GT}}, \text{s.t. } \min_{\mathbf{x}_1 \in \mathbf{P}_{\text{Rec}}} \|\mathbf{x}_1 - \mathbf{x}_2\|_2 < \epsilon \right\} \right|. \end{aligned} \quad (12)$$




Figure 12: Pipeline of QuantDeepSDF.

497 C.2 QuantDeepSDF

498 Compared to DeepSDF, our QuantDeepSDF incorporates the following two modifications:

- 499
- The decoder parameters are quantized to enhance compression efficiency.

500 • To maintain consistency with our NeCGS, the points sampled during training are drawn
501 from TSDF-Def volumes.

502 The pipeline of QuantDeepSDF is shown in Fig. 12. Specifically, the decoder is an MLP, where the
503 input is the concatenated vector of coordinate $\mathbf{x} \in \mathbb{R}^3$ and the i -th embedded feature vector $\mathbf{F}_i \in \mathbb{R}^C$,
504 and the output is the corresponding TSDF-Def value. In our experiment, the decoder consists of 8
505 layers, and the compression ratio is controlled by changing the width of each layer.

506 C.3 Auto-Encoder in Ablation Study

507 Different from the auto-encoder used in our framework, where the embed features are directly
508 optimized, auto-encoder utilizes an encoder to produce the embedded features, where the inputs are
509 the TSDF-Def volumes. And the decoder is kept the same as our framework. During the optimization,
510 the parameters of encoder and decoder are optimized. Once optimized, the embedded features
511 produced by the encoder and decoder parameters are compressed into bitstreams.

512 C.4 More Visual Results

513 Fig. 13 depicts the visual results of the decompressed models from the AMA dataset, DT4D dataset,
514 and Thingi10K dataset under various compression ratios, respectively. With the compression ratio
515 increasing, the decompressed models still preserve the detailed structures, without large distortion.




Figure 13: Visual results of the decompressed models under different compression ratios. From **Top to Bottom**: AMA, DT4D, and Thingi10K. [Q Zoom in for details.](#)

516 **NeurIPS Paper Checklist**

517 **1. Claims**

518 Question: Do the main claims made in the abstract and introduction accurately reflect the
519 paper's contributions and scope?

520 Answer: [Yes]

521 Justification: Abstract.

522 Guidelines:

- 523 • The answer NA means that the abstract and introduction do not include the claims
524 made in the paper.
- 525 • The abstract and/or introduction should clearly state the claims made, including the
526 contributions made in the paper and important assumptions and limitations. A No or
527 NA answer to this question will not be perceived well by the reviewers.
- 528 • The claims made should match theoretical and experimental results, and reflect how
529 much the results can be expected to generalize to other settings.
- 530 • It is fine to include aspirational goals as motivation as long as it is clear that these goals
531 are not attained by the paper.

532 **2. Limitations**

533 Question: Does the paper discuss the limitations of the work performed by the authors?

534 Answer: [Yes]

535 Justification: Sec. 5.

536 Guidelines:

- 537 • The answer NA means that the paper has no limitation while the answer No means that
538 the paper has limitations, but those are not discussed in the paper.
- 539 • The authors are encouraged to create a separate "Limitations" section in their paper.
- 540 • The paper should point out any strong assumptions and how robust the results are to
541 violations of these assumptions (e.g., independence assumptions, noiseless settings,
542 model well-specification, asymptotic approximations only holding locally). The authors
543 should reflect on how these assumptions might be violated in practice and what the
544 implications would be.
- 545 • The authors should reflect on the scope of the claims made, e.g., if the approach was
546 only tested on a few datasets or with a few runs. In general, empirical results often
547 depend on implicit assumptions, which should be articulated.
- 548 • The authors should reflect on the factors that influence the performance of the approach.
549 For example, a facial recognition algorithm may perform poorly when image resolution
550 is low or images are taken in low lighting. Or a speech-to-text system might not be
551 used reliably to provide closed captions for online lectures because it fails to handle
552 technical jargon.
- 553 • The authors should discuss the computational efficiency of the proposed algorithms
554 and how they scale with dataset size.
- 555 • If applicable, the authors should discuss possible limitations of their approach to
556 address problems of privacy and fairness.
- 557 • While the authors might fear that complete honesty about limitations might be used by
558 reviewers as grounds for rejection, a worse outcome might be that reviewers discover
559 limitations that aren't acknowledged in the paper. The authors should use their best
560 judgment and recognize that individual actions in favor of transparency play an impor-
561 tant role in developing norms that preserve the integrity of the community. Reviewers
562 will be specifically instructed to not penalize honesty concerning limitations.

563 **3. Theory Assumptions and Proofs**

564 Question: For each theoretical result, does the paper provide the full set of assumptions and
565 a complete (and correct) proof?

566 Answer: [NA]

567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Sec. 4.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672

Answer: [Yes]

Justification: We include the code in the supplemental material.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Sec. 4.1

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Sec. 4.2 and 4.3.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- 673 • It is OK to report 1-sigma error bars, but one should state it. The authors should
674 preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis
675 of Normality of errors is not verified.
- 676 • For asymmetric distributions, the authors should be careful not to show in tables or
677 figures symmetric error bars that would yield results that are out of range (e.g. negative
678 error rates).
- 679 • If error bars are reported in tables or plots, The authors should explain in the text how
680 they were calculated and reference the corresponding figures or tables in the text.

681 8. Experiments Compute Resources

682 Question: For each experiment, does the paper provide sufficient information on the com-
683 puter resources (type of compute workers, memory, time of execution) needed to reproduce
684 the experiments?

685 Answer: [Yes]

686 Justification: Sec. 4.1 and 4.3.

687 Guidelines:

- 688 • The answer NA means that the paper does not include experiments.
- 689 • The paper should indicate the type of compute workers CPU or GPU, internal cluster,
690 or cloud provider, including relevant memory and storage.
- 691 • The paper should provide the amount of compute required for each of the individual
692 experimental runs as well as estimate the total compute.
- 693 • The paper should disclose whether the full research project required more compute
694 than the experiments reported in the paper (e.g., preliminary or failed experiments that
695 didn't make it into the paper).

696 9. Code Of Ethics

697 Question: Does the research conducted in the paper conform, in every respect, with the
698 NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

699 Answer: [Yes]

700 Justification: [NA]

701 Guidelines:

- 702 • The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- 703 • If the authors answer No, they should explain the special circumstances that require a
704 deviation from the Code of Ethics.
- 705 • The authors should make sure to preserve anonymity (e.g., if there is a special consid-
706 eration due to laws or regulations in their jurisdiction).

707 10. Broader Impacts

708 Question: Does the paper discuss both potential positive societal impacts and negative
709 societal impacts of the work performed?

710 Answer: [Yes]

711 Justification: [NA]

712 Guidelines:

- 713 • The answer NA means that there is no societal impact of the work performed.
- 714 • If the authors answer NA or No, they should explain why their work has no societal
715 impact or why the paper does not address societal impact.
- 716 • Examples of negative societal impacts include potential malicious or unintended uses
717 (e.g., disinformation, generating fake profiles, surveillance), fairness considerations
718 (e.g., deployment of technologies that could make decisions that unfairly impact specific
719 groups), privacy considerations, and security considerations.
- 720 • The conference expects that many papers will be foundational research and not tied
721 to particular applications, let alone deployments. However, if there is a direct path to
722 any negative applications, the authors should point it out. For example, it is legitimate
723 to point out that an improvement in the quality of generative models could be used to

724 generate deepfakes for disinformation. On the other hand, it is not needed to point out
725 that a generic algorithm for optimizing neural networks could enable people to train
726 models that generate Deepfakes faster.

- 727 • The authors should consider possible harms that could arise when the technology is
728 being used as intended and functioning correctly, harms that could arise when the
729 technology is being used as intended but gives incorrect results, and harms following
730 from (intentional or unintentional) misuse of the technology.
- 731 • If there are negative societal impacts, the authors could also discuss possible mitigation
732 strategies (e.g., gated release of models, providing defenses in addition to attacks,
733 mechanisms for monitoring misuse, mechanisms to monitor how a system learns from
734 feedback over time, improving the efficiency and accessibility of ML).

735 11. Safeguards

736 Question: Does the paper describe safeguards that have been put in place for responsible
737 release of data or models that have a high risk for misuse (e.g., pretrained language models,
738 image generators, or scraped datasets)?

739 Answer: [NA]

740 Justification: [NA]

741 Guidelines:

- 742 • The answer NA means that the paper poses no such risks.
- 743 • Released models that have a high risk for misuse or dual-use should be released with
744 necessary safeguards to allow for controlled use of the model, for example by requiring
745 that users adhere to usage guidelines or restrictions to access the model or implementing
746 safety filters.
- 747 • Datasets that have been scraped from the Internet could pose safety risks. The authors
748 should describe how they avoided releasing unsafe images.
- 749 • We recognize that providing effective safeguards is challenging, and many papers do
750 not require this, but we encourage authors to take this into account and make a best
751 faith effort.

752 12. Licenses for existing assets

753 Question: Are the creators or original owners of assets (e.g., code, data, models), used in
754 the paper, properly credited and are the license and terms of use explicitly mentioned and
755 properly respected?

756 Answer: [Yes]

757 Justification: [NA]

758 Guidelines:

- 759 • The answer NA means that the paper does not use existing assets.
- 760 • The authors should cite the original paper that produced the code package or dataset.
- 761 • The authors should state which version of the asset is used and, if possible, include a
762 URL.
- 763 • The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- 764 • For scraped data from a particular source (e.g., website), the copyright and terms of
765 service of that source should be provided.
- 766 • If assets are released, the license, copyright information, and terms of use in the
767 package should be provided. For popular datasets, paperswithcode.com/datasets
768 has curated licenses for some datasets. Their licensing guide can help determine the
769 license of a dataset.
- 770 • For existing datasets that are re-packaged, both the original license and the license of
771 the derived asset (if it has changed) should be provided.
- 772 • If this information is not available online, the authors are encouraged to reach out to
773 the asset's creators.

774 13. New Assets

775 Question: Are new assets introduced in the paper well documented and is the documentation
776 provided alongside the assets?

777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via regular templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.