

Geometry Distributions

Supplementary Material




Figure 15. The ground-truth geometries for the shapes shown in Fig. 1.

name	mesh stats		
	# vtx (K)	# face (K)	disk storage (Mb)
WuKong	19226	6408	878
Archimedes	6162	2054	276
loong	3956	1318	176
jellyfish	3908	1302	178
mouse	793	1423	58
archer	710	1420	58
lamp	478	159	21
warrior	433	866	35
lion	303	606	24
jacket	146	49	11
city	249	83	11
Parthenon	117	39	5
valley	74	37	6
Spot	3.2	5.8	0.37
table	0.82	0.73	0.07

Table A.1. We report the mesh complexity and storage cost for the experimented shapes. Unless explicitly mentioned in the table, the shapes can be found in Fig. 15.

A. Additional results

Fig. 16 extends Fig. 2, demonstrating that our approach generates more uniformly distributed samples with higher fidelity across different resolutions, compared to the vector field-based method. Fig. 18 (an extension of Tab. 3c), and Fig. 17 (an extension of Fig. 13), provide error




Figure 16. Additional results on comparison to vector fields-based method (extension of Fig. 2), our **GEOMDIST** produces more uniformly distributed samples with higher fidelity. *Top*: vector fields. *Bottom*: ours. From left to right, we show results of 10 thousands, 100 thousands and 1 million points.

visualization of L_2 distance to surface, demonstrating how the sampling steps affect the forward sampling and inverse sampling, respectively, for geometry recovery. In Fig. 19 and Fig. 20, we show additional results of using our method to represent textured geometry and high-resolution scene. In Fig. 21 we show additional re-

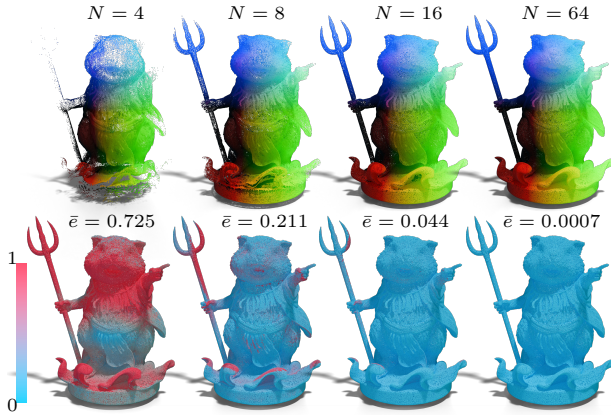


Figure 17. *Top*: the recovered shape from inversion $\mathcal{E} \circ \mathcal{D}(\mathbf{x})$ using different number of inverse sampling steps as in Eq. (4). *Bottom*: we color the per-sample error of the inversion on the original shape and report the average error $\bar{\epsilon}$ (L_2 distance to the surface).




Figure 18. Errors (L_2 distance to the surface) of 1 million points using different number of sampling steps N as in Eq. (3). *Left*: $N = 8$. *Middle*: $N = 64$. *Right*: ground-truth mesh. The per-point color is visualized according to the error (L_2 distance to the ground-truth surface). The results of $N = 16$ and $N = 32$ are very close as shown in Tab. 3c, thus are omitted here.

sults to justify the inversion process discussed in Sec. 3.2 and Sec. 4.4: we composite inverse sampling and forward sampling from *different* surfaces, yet still obtain expected results. Specifically, the composed sampling process allows us to transform the WuKong shape into a jellyfish, lamp, sphere, or a plane. In Tab. A.1 we report the mesh complexity of the experimented shapes.

B. Implementation details

B.1. Mesh normalization

We normalize all the meshes using the following pseudocode. First, we sample 10 million points on the surface. Next, we shift and scale the mesh based on the mean and standard deviation of the sampled points. As a result, the surface points are approximately centered around zero with unit variance. We found that this normalization is effective in stabilizing the training process.



Figure 19. Applications on textured geometries. The setup is the same with Fig. 9. *Top*: ground-truth. *Bottom*: ours.




Figure 20. Results on city. *Top*: ground-truth. *Bottom*: ours.

```

1 points, _ = trimesh.sample.sample_surface(mesh,
2 mesh.vertices -= points.mean()

```

```
3 mesh.vertices /= points.std()
```

B.2. Uniform distribution

When using uniform distribution as the initial noise in diffusion models (e.g., see Fig. 4), we scale the samples from uniform distribution to have zero mean and unit variance.

```
1 n = (torch.rand_like(x) - 0.5) / np.sqrt(1/12)
```

B.3. Chamfer distance

We use Chamfer distance to measure the distance between samples from our Geometry distribution, \mathcal{X}_{gen} , and the samples from ground-truth surface, \mathcal{X}_{ref} , to quantify the accuracy. This is defined as:

$$\text{ChamferDist}(\mathcal{X}_{\text{ref}}, \mathcal{X}_{\text{gen}}) = \frac{1}{|\mathcal{X}_{\text{ref}}|} \sum_{\mathbf{a} \in \mathcal{X}_{\text{ref}}} \min_{\mathbf{b} \in \mathcal{X}_{\text{gen}}} \|\mathbf{a} - \mathbf{b}\|_2 + \frac{1}{|\mathcal{X}_{\text{gen}}|} \sum_{\mathbf{b} \in \mathcal{X}_{\text{gen}}} \min_{\mathbf{a} \in \mathcal{X}_{\text{ref}}} \|\mathbf{a} - \mathbf{b}\|_2.$$

The python code for calculating the Chamfer distance is as follows:

```
1 # prediction: B x 3
2 # reference: B x 3
3 from scipy.spatial import cKDTree as KDTree
4 tree = KDTree(prediction)
5 dist, _ = tree.query(reference)
6 d1 = dist
7 gt_to_gen_chamfer = np.mean(dist)
8 gt_to_gen_chamfer_sq = np.mean(np.square(dist))
9
10 tree = KDTree(reference)
11 dist, _ = tree.query(prediction)
12 d2 = dist
13 gen_to_gt_chamfer = np.mean(dist)
14 gen_to_gt_chamfer_sq = np.mean(np.square(dist))
15
16 cd = gt_to_gen_chamfer + gen_to_gt_chamfer
```

B.4. Sampling algorithm

We show the sampling algorithm (Algorithm 2) proposed in EDM [19] for completeness.

Algorithm 2 Sampling

```
1: procedure SAMPLING( $\mathbf{x}, t_i \in \{0, \dots, N\}$ )
2:    $\mathbf{x}_0 = t_0 \mathbf{n}$  where  $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$ 
3:   for  $i \in \{0, 1, \dots, N-1\}$  do
4:      $\mathbf{d}_i = (\mathbf{x}_i - D_\theta(\mathbf{x}_i, t_i)) / t_i$ 
5:      $\mathbf{x}_{i+1} = \mathbf{x}_i + (t_{i+1} - t_i) \cdot \mathbf{d}_i$ 
6:   end for
7: end procedure
```

B.5. Networks

The forward passes of the middle blocks and final block (illustrated in Fig. 6) are implemented as follows:

```
1 # x: B x C
2 # t: B x C
3 # middle block
4 c = emb_mp_linear(t, gain=emb_gain) + 1
5 x = normalize(x)
6 res = x_pre_mp_linear(mp_silu(x))
7 res = mp_silu(res * c.to(y.dtype))
8 res = x_post_mp_linear(res)
9 x = mp_sum(x, res, t=0.3)
```

```
1 # x: B x C
2 # t: B x C
3 # final block
4 c = emb_mp_linear(t, gain=final_emb_gain) + 1
5 x = x_pre_mp_linear(mp_silu(normalize(x)))
6 x = mp_silu(x * c.to(y.dtype))
7 out = x_post_mp_linear(x, gain=final_out_gain)
```

B.6. Vector fields

The vector fields are coordinate-based networks which outputs vectors pointing towards to the surface. We use Libigl library [17] to process the data.

```
1 # p: B x 3
2 p = np.random.randn(B, 3)
3 v, f = igl.read_triangle_mesh(obj_path)
4 d, _, c = igl.point_mesh_squared_distance(p, v, f)
5 unsigned_distances = np.sqrt(d) # B
6 vectors = c - p # B x 3
```

B.7. Color fields

We use the hashing grids proposed by Instant-NGP to implement the color field network in Fig. 10. The implementation is from the official github repository.

```
1 encoder_config = """{
2   "otype": "HashGrid",
3   "n_levels": 16,
4   "n_features_per_level": 2,
5   "log2_hashmap_size": 19,
6   "base_resolution": 16,
7 }"""
8
9
10 network_config = """{
11   "otype": "FullyFusedMLP",
12   "activation": "ReLU",
13   "output_activation": "Sigmoid",
14   "n_neurons": 64,
15   "n_hidden_layers": 2
16 }"""
17
18 class ColorField(nn.Module):
19     def __init__(self):
20         super().__init__()
21         self.encoding = tcnn.Encoding(
22             n_input_dims=3, encoding_config=json.loads(
23                 encoder_config))
24         self.network = tcnn.Network(
25             n_input_dims=self.encoding.n_output_dims,
26             n_output_dims=3, network_config=json.loads(
27                 network_config))
```

```
24 def forward(self, x):
25     x = self.encoding(x)
26     x = self.network(x)
27     return x
```

We optimize L_1 -loss between the predicted and ground-truth colors. The training takes around 3 minutes. When the training is done, we can query colors for all spatial points,

$$\text{ColorField}(\mathbf{x}) = \mathbf{c}. \quad (\text{B.1})$$

The colors in Fig. 10 are obtained by $\text{ColorField}(\mathcal{E}(\mathbf{n}))$ where $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$.

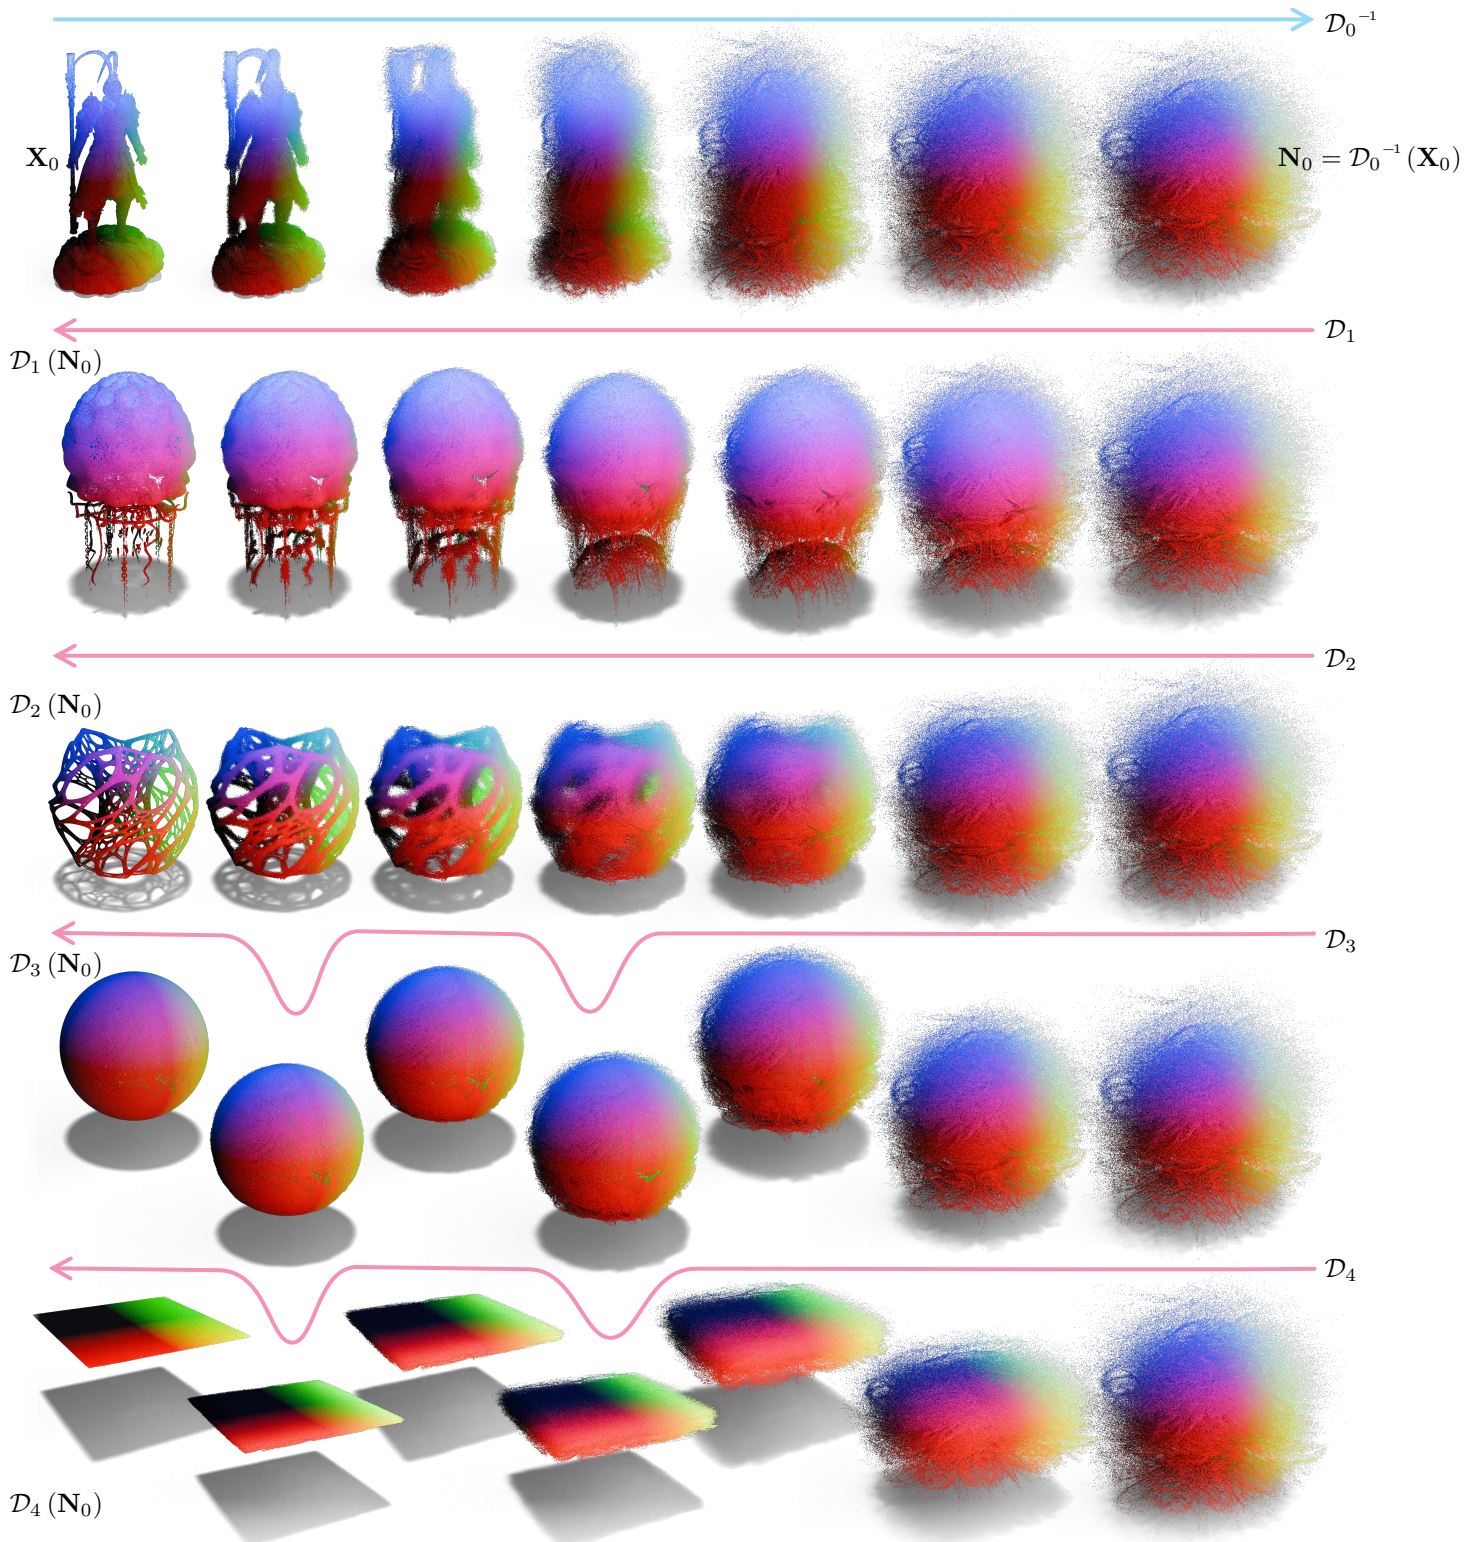


Figure 21. We denote the trained diffusion networks that map from a Gaussian distribution to the Wukong, jellyfish, lamp, sphere, and plane mesh as $\mathcal{D}_0, \mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3, \mathcal{D}_4$, respectively. We sample 1 million points from the WuKong mesh, denoting these samples as \mathbf{X}_0 . Applying the inversion we obtain $\mathbf{N}_0 = \mathcal{D}_0^{-1}(\mathbf{X}_0)$. In the leftmost column, we show the samples $\mathcal{D}_1(\mathbf{N}_0), \mathcal{D}_2(\mathbf{N}_0), \mathcal{D}_3(\mathbf{N}_0), \mathcal{D}_4(\mathbf{N}_0)$, which closely approximate the original shapes, demonstrating the accuracy of our trained diffusion nets. For \mathcal{D}_0^{-1} , we show results at timesteps 0, 10, 15, 20, 25, 30, 64 aligned from left to right. For $\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3, \mathcal{D}_4$ we show results at timesteps 30, 40, 45, 48, 52, 55, 64, aligned from right to left.