Discrete Point Flow Networks for Efficient Point Cloud Generation Supplementary Material

Roman Klokov[1][0000−0001−9592−7009], Edmond Boyer[1][0000−0002−1182−3729], and Jakob Verbeek[2][0000−0003−1419−1816]

[1] Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France
{firsname.lastname}@inria.fr
[2] Facebook AI Research

In this supplementary material we present training details in Section 1. In Section 2 we repeat generative task evaluation results, this time including the standard deviations for all the metrics. We provide additional samples for qualitative comparison of our DPF-Nets with l-GANs, AtlasNet, and PointFlow in the sparse and dense autoencoding setups in Section 3. In Section 4 we show multiple instances of generation for DPF-Nets and PointFlow. Further we show additional latent shape space interpolation examples, and visualize point cloud evolution across the flow for generative class specific models in sections 5 and 6.

1 Training Details

All the models were trained with AMSGrad optimizer [3] with decoupled weight decay regularization [2] with step-like schedule for the learning rate. All the hyperparameters for all the experiments can be found on the paper webpage https://github.com/Regenerator/dpf-nets.

2 Detailed Generative Modeling Performance

In Table 1 we present the same evaluation as in Table 2 of the main paper, but include the standard deviations across the ten sets of point clouds sampled from each model.

3 Qualitative Results for Autoencoding

In Figure 1 and Figure 2 we show reconstruction results for l-GANs, AtlasNet, PointFlow, and ours DPF-Net from sparse and dense input point clouds. All the models used 512 and 32,768 points as a sparse or dense input accordingly. Note, that AtlasNet, PointFlow and DPF-Net are able to reconstruct arbitrary densely (32,768 points here), for l-GANs output size is fixed to 2,048.

4 Qualitative Results for Generation

In figures 3—5 we provide additional samples from the models trained for the Airplane, Car, and Chair classes. For each point cloud we sample $32,768$ points.

5 Additional interpolations

In figures 6—8 we provide latent space interpolations from the models trained for the Airplane, Car, and Chair classes. We sample two shapes from the test data (left- and right-most in the figures), and interpolate between corresponding latent variable samples to obtain a path in the shape space. For each latent shape on the path we then sample a point cloud of size $32,768$.

6 Additional flow illustrations

In figures 9—11 we provide examples of the generating flow for sampled point clouds from the classes Airplane, Car, and Chair. We sample a shape from data, obtain initial Gaussian from the corresponding latent variable, sample $32,768$ points by the flow, and then visualize the evolution of the point cloud across the initial Gaussian and layers 32, 48, 56, 60, 62, 63 of the generative flow network.

Table 1: Generative modeling evaluation. JSD and MMD-EMD are multiplied by 10^2, MMD-CD by 10^4. Cases where the oracle does not obtain best results are underlined.

Category	Model	JSD↓	MMD↓		COV↑, %		1-NNA↓, %	
			CD	EMD	CD	EMD	CD	EMD
Airplane	1-GAN-CD [1]	2.76 ± 0.16	**5.69** ± 0.04	5.16 ± 0.02	39.5 ± 0.8	17.1 ± 0.6	72.9 ± 0.8	92.1 ± 0.6
	1-GAN-EMD [1]	1.77 ± 0.13	6.05 ± 0.04	**4.15** ± 0.02	39.7 ± 1.4	40.4 ± 1.2	75.7 ± 0.6	73.0 ± 1.2
	PointFlow [4]	1.42 ± 0.12	6.05 ± 0.05	4.32 ± 0.01	**44.7** ± 1.2	**48.4** ± 1.0	**70.9** ± 1.0	**68.4** ± 1.0
	DPF-Nets (Ours)	**0.94** ± 0.11	6.07 ± 0.04	4.26 ± 0.02	**46.8** ± 1.2	**48.4** ± 0.9	**70.6** ± 1.0	**67.0** ± 1.2
	Oracle	0.50 ± 0.04	5.97 ± 0.09	3.98 ± 0.01	51.4 ± 1.0	52.7 ± 1.3	49.8 ± 1.3	48.2 ± 1.1
Car	1-GAN-CD [1]	2.65 ± 0.07	**8.83** ± 0.06	5.36 ± 0.01	41.3 ± 0.8	15.9 ± 1.3	**62.6** ± 0.6	92.7 ± 0.4
	1-GAN-EMD [1]	1.31 ± 0.10	**9.00** ± 0.08	**4.40** ± 0.01	38.3 ± 1.2	32.9 ± 0.7	65.2 ± 0.4	**63.2** ± 1.0
	PointFlow [4]	0.59 ± 0.02	9.53 ± 0.06	4.71 ± 0.01	**42.3** ± 1.0	35.8 ± 1.3	70.1 ± 0.9	74.2 ± 0.6
	DPF-Nets (Ours)	**0.45** ± 0.02	9.59 ± 0.04	4.61 ± 0.01	**43.4** ± 0.9	**45.8** ± 1.0	70.3 ± 0.6	**64.3** ± 1.5
	Oracle	0.37 ± 0.03	9.24 ± 0.06	4.56 ± 0.01	52.8 ± 1.1	52.7 ± 0.9	50.9 ± 1.1	50.5 ± 1.2
Chair	1-GAN-CD [1]	3.65 ± 0.09	**16.66** ± 0.08	7.91 ± 0.02	42.3 ± 0.5	17.1 ± 0.5	68.5 ± 0.5	96.5 ± 0.1
	1-GAN-EMD [1]	1.27 ± 0.06	**16.78** ± 0.07	**5.75** ± 0.01	44.3 ± 0.9	43.8 ± 1.0	66.6 ± 0.6	67.8 ± 0.7
	PointFlow [4]	1.51 ± 0.11	17.15 ± 0.10	6.20 ± 0.01	43.3 ± 0.8	**46.5** ± 1.0	67.0 ± 0.3	70.4 ± 0.6
	DPF-Nets (Ours)	**1.01** ± 0.06	17.08 ± 0.11	6.14 ± 0.01	**46.9** ± 0.8	**48.5** ± 1.1	**63.5** ± 1.3	**64.8** ± 0.7
	Oracle	0.49 ± 0.01	16.39 ± 0.07	5.71 ± 0.01	52.8 ± 0.8	53.4 ± 1.1	49.7 ± 0.7	49.6 ± 0.9

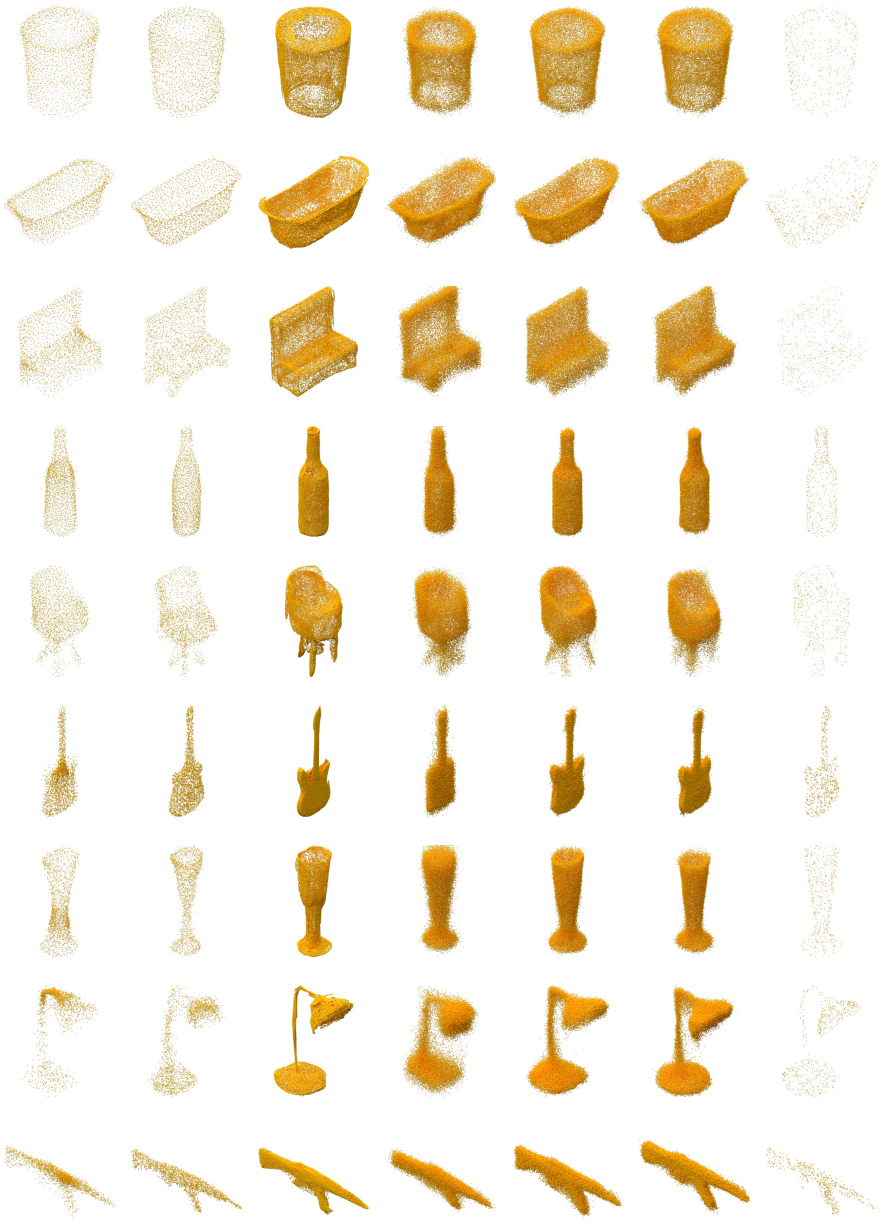

Fig. 1: Qualitative comparison in the autoencoding task with sparse inputs. Left to right: reconstructions from l-GAN-CD, l-GAN-EMD, AtlasNet, PointFlow, DPF-Nets (orig.), DPF-Nets (norm.), and ground-truth.

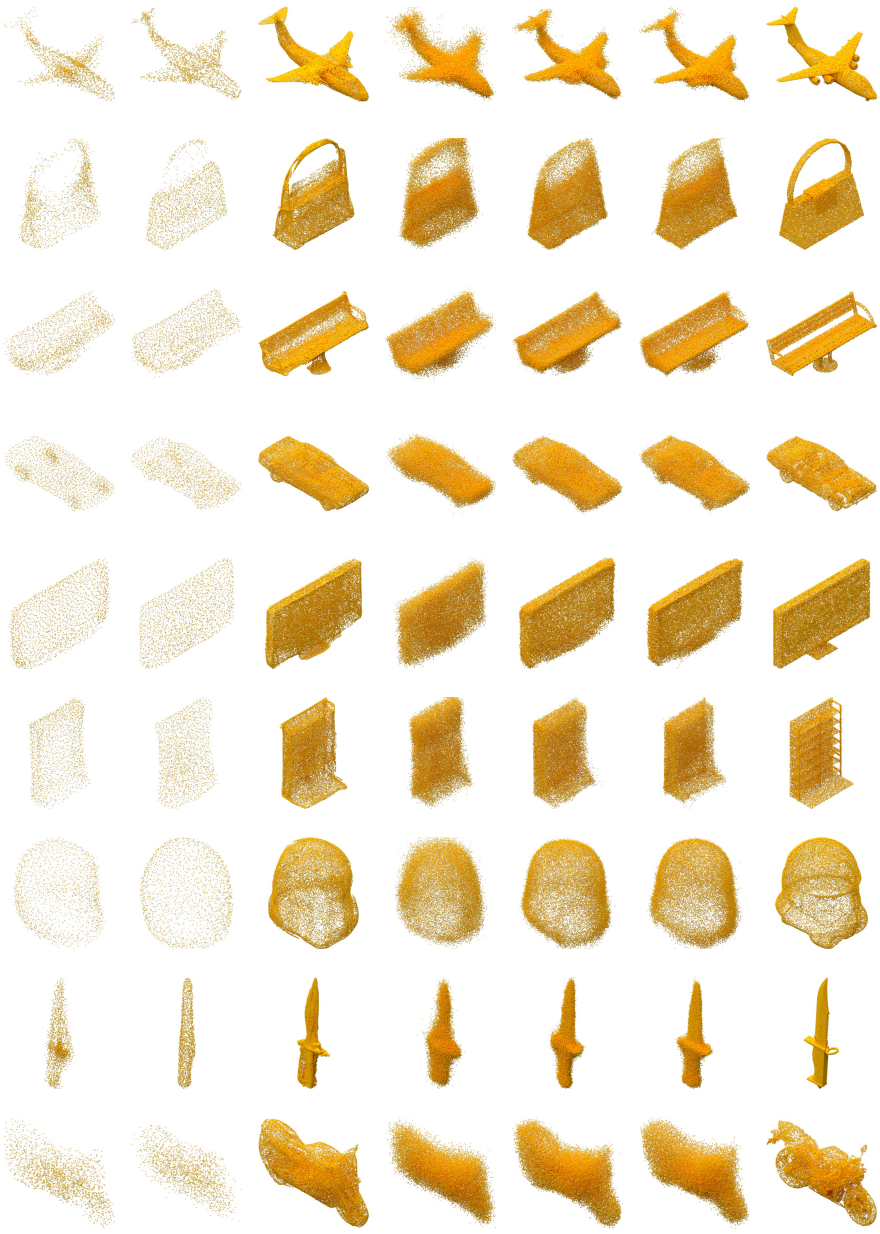

Fig. 2: Qualitative comparison in the autoencoding task with dense inputs. Left to right: reconstructions from l-GAN-CD, l-GAN-EMD, AtlasNet, PointFlow, DPF-Nets (orig.), DPF-Nets (norm.), and ground-truth.

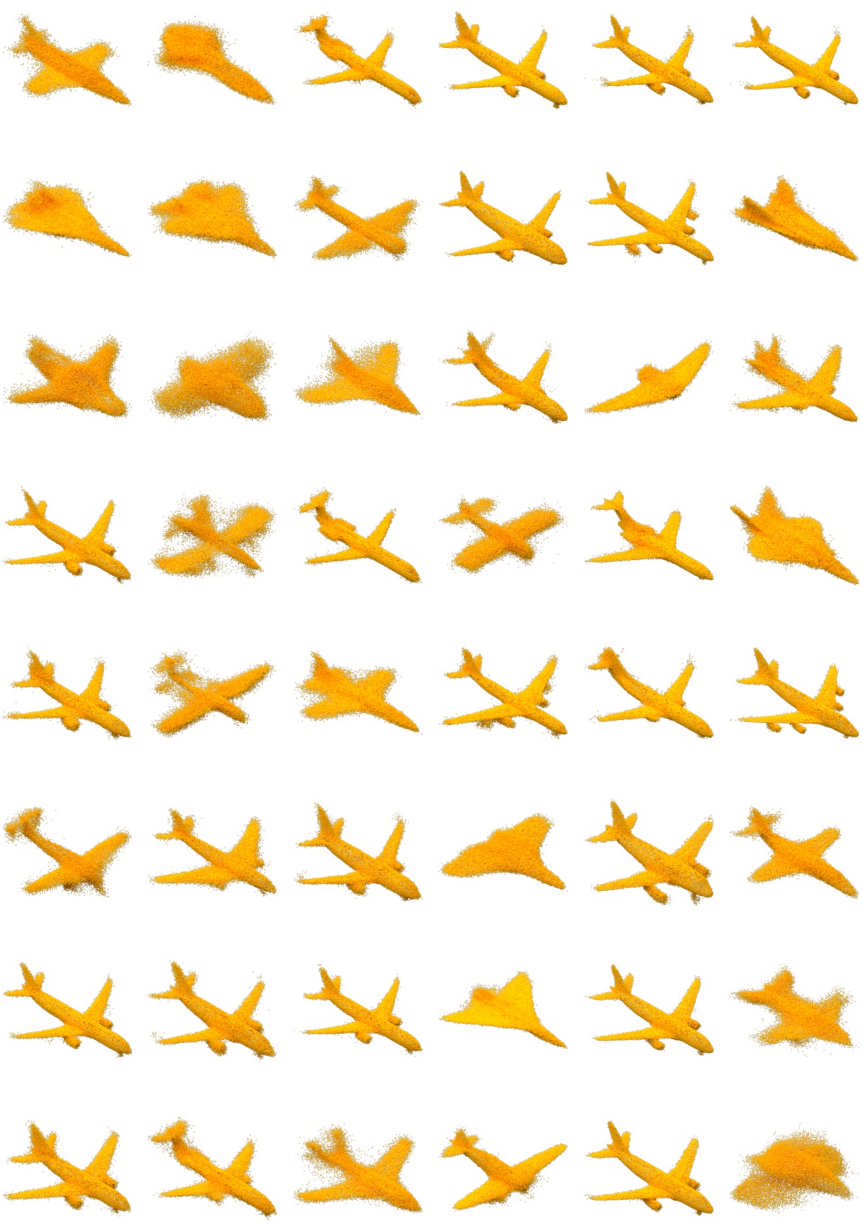

Fig. 3: Random samples from models trained on the Airplane class. Columns 1—3 samples from PointFlow, columns 4—6 samples from DPF-Net.

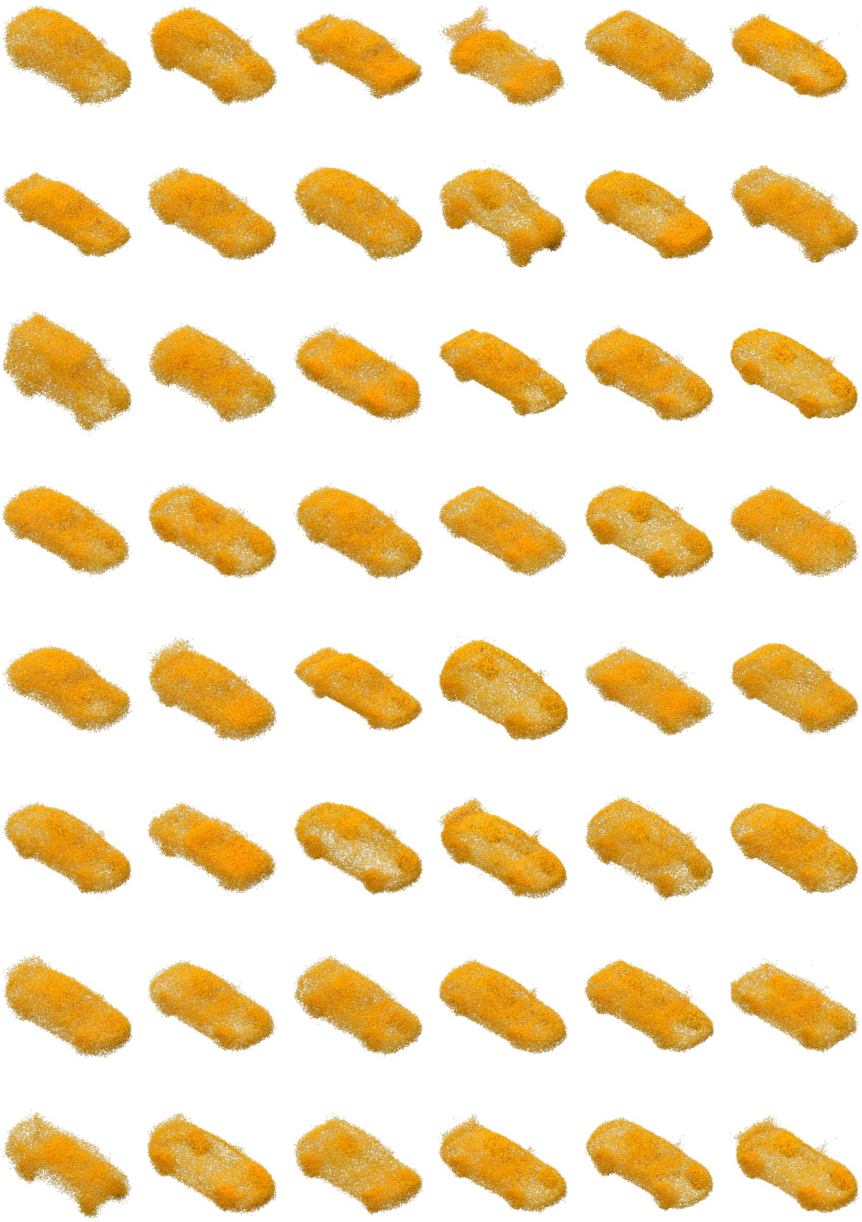

Fig. 4: Random samples from models trained on the Car class. Columns 1—3 samples from PointFlow, columns 4—6 samples from DPF-Net.

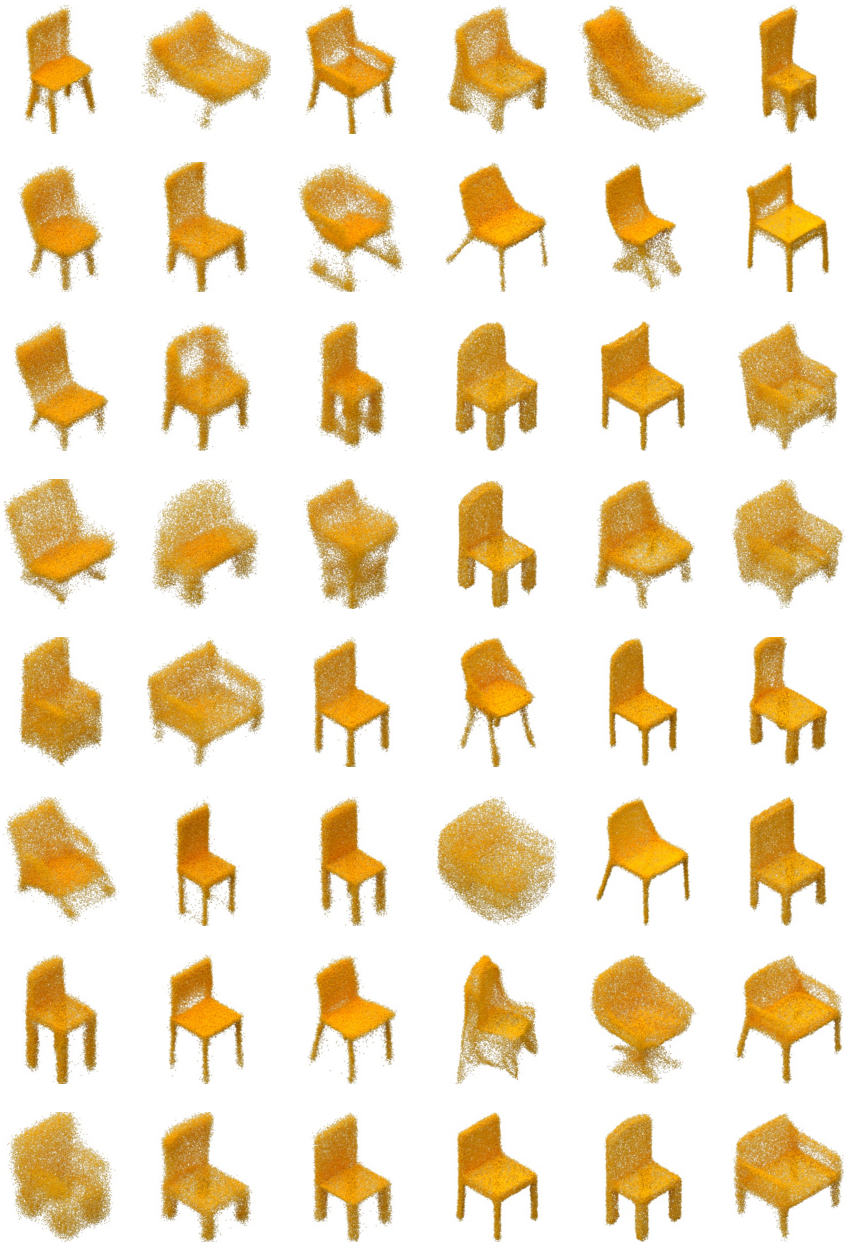

Fig. 5: Random samples from models trained on the Chair class. Columns 1—3 samples from PointFlow, columns 4—6 samples from DPF-Net.

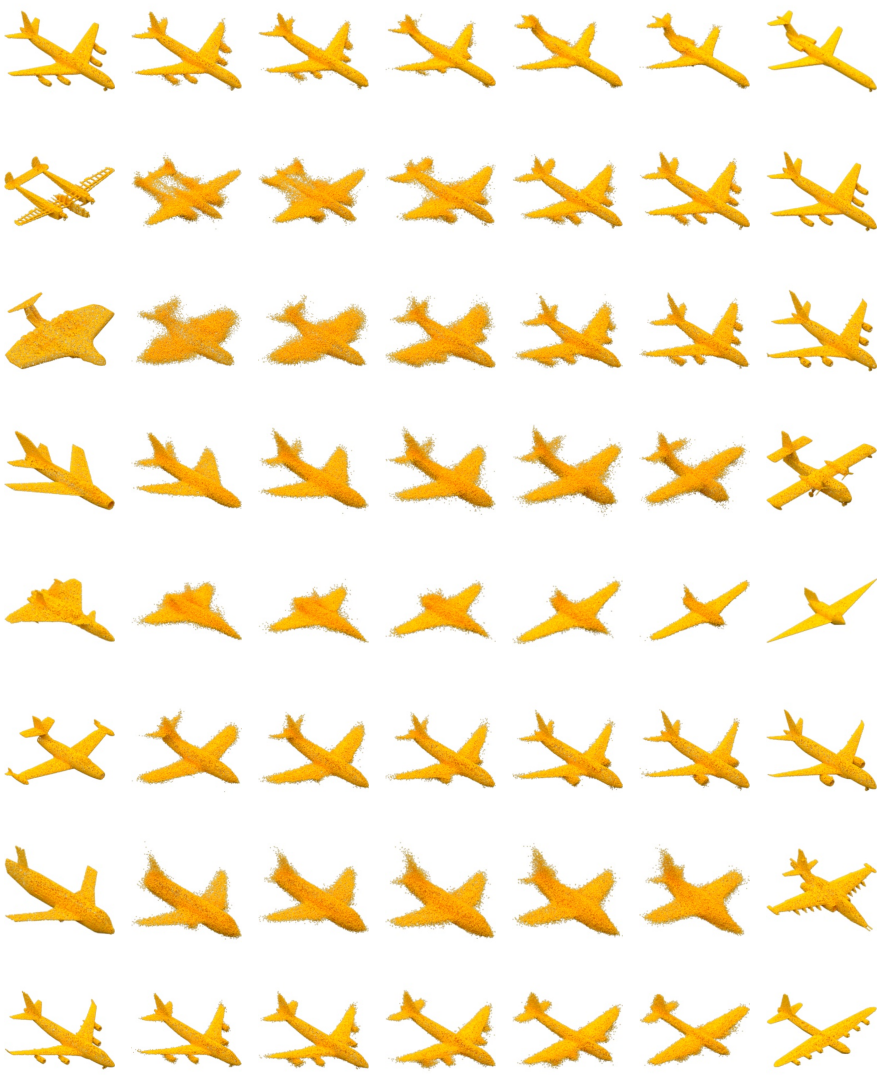

Fig. 6: Interpolations between Airplane data samples.

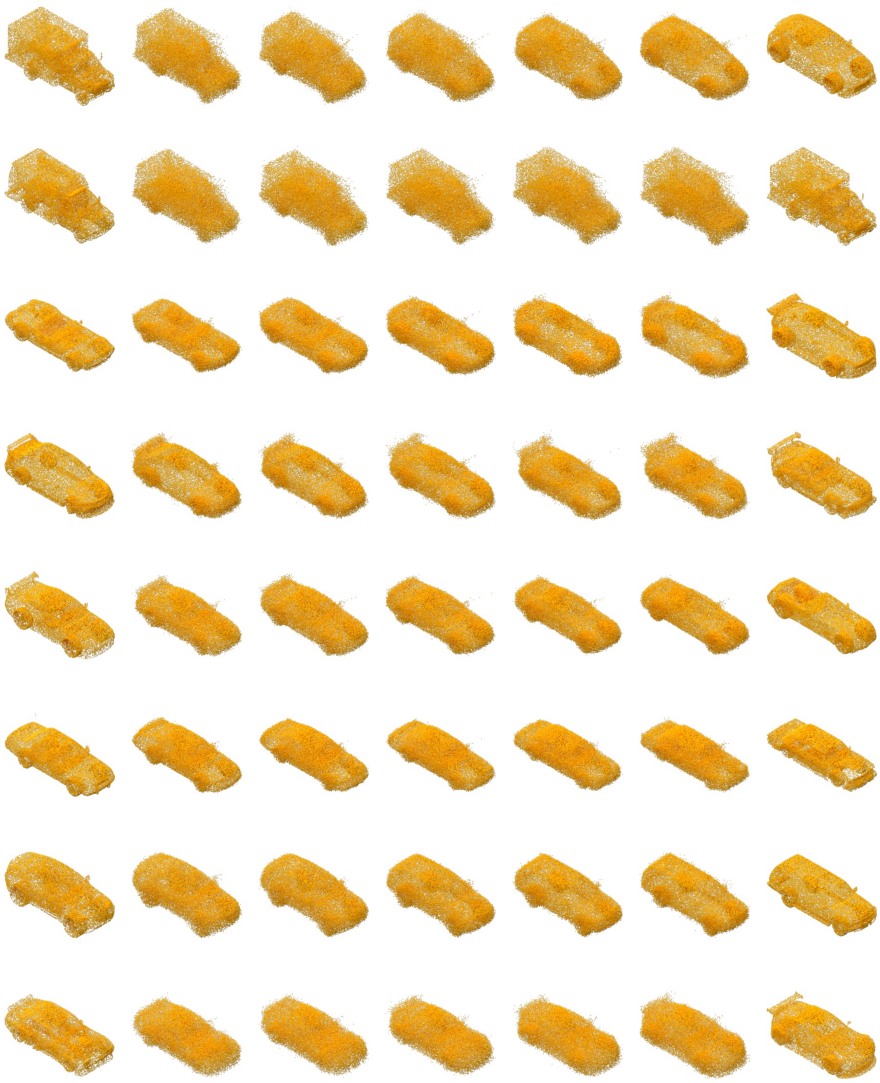

Fig. 7: Interpolations between Car data samples.

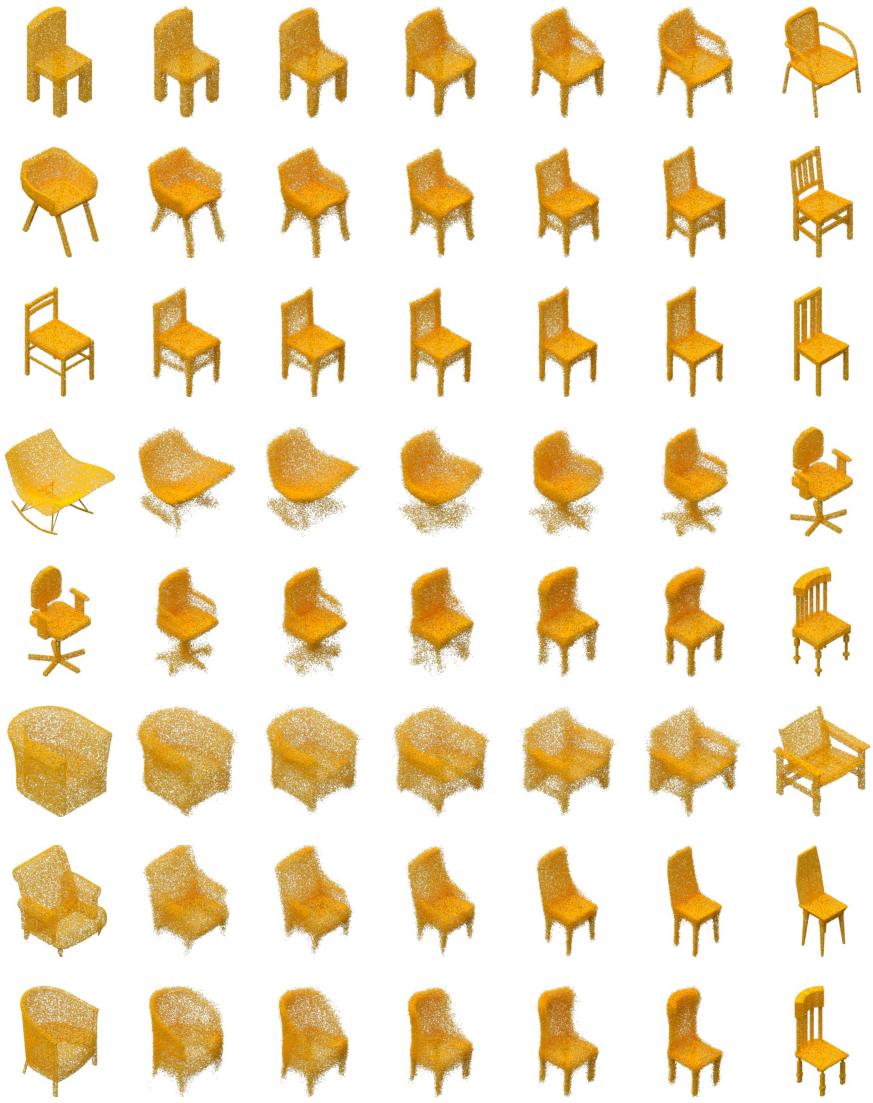

Fig. 8: Interpolations between Chairs data samples.

Fig. 9: Generating flow for sampled point clouds from the Airplane model.

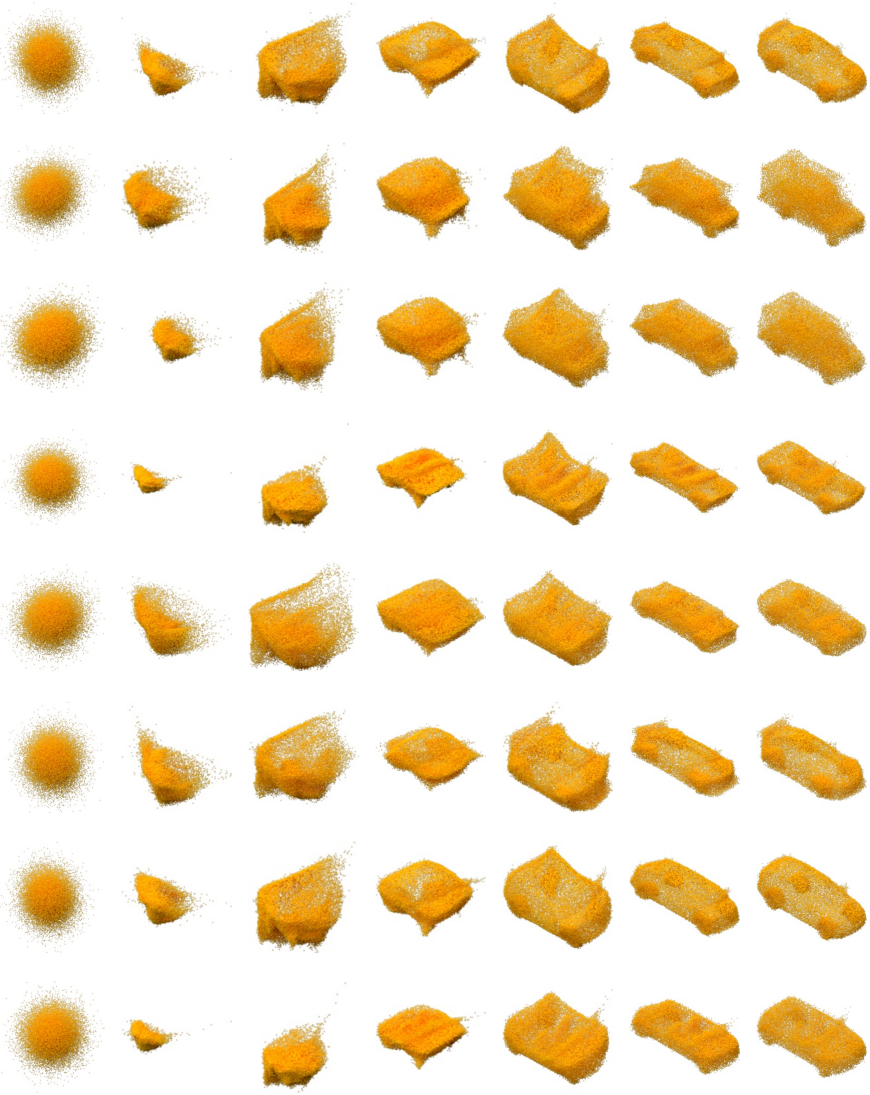

Fig. 10: Generating flow for sampled point clouds from the Car model.

Fig. 11: Generating flow for sampled point clouds from the Chair model.

References

1. Achlioptas, P., Diamanti, O., Mitliagkas, I., Guibas, L.: Learning representations and generative models for 3D point clouds. In: ICML (2018)
2. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: ICLR (2019)
3. Reddi, S., Kale, S., Kumar, S.: On the convergence of Adam and beyond. In: ICLR (2018)
4. Yang, G., Huang, X., Hao, Z., Liu, M.Y., Belongie, S., Hariharan, B.: PointFlow: 3D point cloud generation with continuous normalizing flows. In: ICCV (2019)