

Mosaic-SDF for 3D Generative Models

Lior Yariv^{1,3} *

¹GenAI, Meta

Omri Puny³

²FAIR, Meta

Oran Gafni¹

³Weizmann Institute of Science

Yaron Lipman^{2,3}

Abstract

Current diffusion or flow-based generative models for 3D shapes divide to two: distilling pre-trained 2D image diffusion models, and training directly on 3D shapes. When training a diffusion or flow models on 3D shapes a crucial design choice is the shape representation. An effective shape representation needs to adhere three design principles: it should allow an efficient conversion of large 3D datasets to the representation form; it should provide a good tradeoff of approximation power versus number of parameters; and it should have a simple tensorial form that is compatible with existing powerful neural architectures. While standard 3D shape representations such as volumetric grids and point clouds do not adhere to all these principles simultaneously, we advocate in this paper a new representation that does. We introduce Mosaic-SDF (M-SDF): a simple 3D shape representation that approximates the Signed Distance Function (SDF) of a given shape by using a set of local grids spread near the shape’s boundary. The M-SDF representation is fast to compute for each shape individually making it readily parallelizable; it is parameter efficient as it only covers the space around the shape’s boundary; and it has a simple matrix form, compatible with Transformer-based architectures. We demonstrate the efficacy of the M-SDF representation by using it to train a 3D generative flow model including class-conditioned generation with the ShapeNetCore-V2 (3D Warehouse) dataset, and text-to-3D generation using a dataset of about 600k caption-shape pairs.

1. Introduction

Image-based generative models have rapidly advanced in recent years due to improvements in generation methodologies (e.g., Diffusion Models), model architecture and conditioning (e.g., Text-to-Image Models, Attention/Transformer layers), and the consolidation of large, high quality image datasets. Although generation of 3D shapes have progressed as well, it has not seen the same level of progress demonstrated in image generation.

*Work done while interning at Meta.

Current works for 3D generation divide mostly into two groups: *Optimization based*: 2D image generative priors are used for training/optimizing 3D shape [20, 35, 46]. The main benefit is leveraging existing powerful image models, but the generation process is expensive/slow as it requires training a new model for each sample. Furthermore, using only image priors is often insufficient to build consistent 3D shape (e.g., the Janus effect) [7]. *Forward based*: the 3D shapes are generated by a forward process, such as solving an Ordinary Differential Equation (ODE), making generation process more efficient than optimization based methods. However, forward based model are first trained on a dataset of 3D shapes using, e.g., Diffusion or Flow Models, where to that end shapes are first transformed into some canonical 3D representation, e.g., volumetric grid or a point cloud. Forward based models work directly with 3D shapes and suffer from two limitations: first, they require 3D shape datasets for training and these still lag behind image datasets in terms of quantity and quality. Second, in contrast to images, 3D shapes do not occupy the full 3D space and therefore an effective 3D shape representation is more challenging to find/work with.

In this work we focus on forward based 3D generative models. To enable high quality, large scale 3D generation the 3D shape representation of choice should adhere the following design principles:

- (i) **Preprocess efficiency**: can be efficiently computed for a large collection of shapes.
- (ii) **Parameter efficiency**: provides a good approximation vs. parameter count trade-off.
- (iii) **Simple structure**: has a simple tensorial structure, compatible with expressive neural architectures.

Examining existing 3D representations used for training generative 3D models including volumetric grids [19, 31], tri-planes [10, 41, 45], point clouds [26, 33, 48, 52], meshes [23] and neural fields [15] we find that all are lacking in one or more of the above design principles. For example, volumetric grids scale cubically and maintain redundant information in the entire space, while tri-planes tend to be more efficient but still encode full 3D tensors and require a neural network trained jointly on the training set. Point clouds do not provide a full surface information; meshes do not enjoy

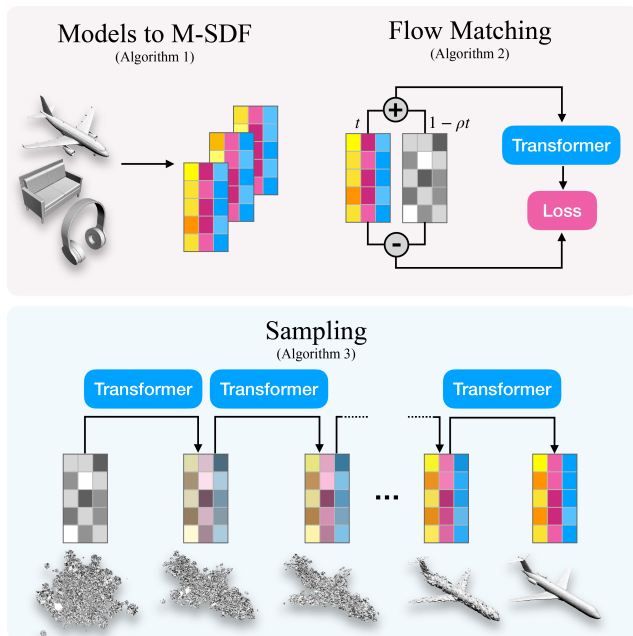


Figure 1. Method overview. Train (top): First we convert the dataset of shapes to M-SDF representations (Algorithm 1), next we train a Flow Matching model with the M-SDF representations (Algorithm 2). Sampling (bottom): We random a noisy M-SDF and numerically solve the ode in equation 13 (Algorithm 3).

a simple tensorial structure in general; and neural fields are encoded by weight vector with non-trivial symmetries [32].

The goal of this work is to introduce Mosaic-SDF (M-SDF), a simple novel 3D shape representation that achieves all desired properties listed above. With M-SDF we are able to train a forward based flow generative model on a datasets of 3D shapes achieving high quality and diversity. In a nutshell, M-SDF approximates an arbitrary Signed Distance Function (SDF) with a *set* of small ($7 \times 7 \times 7$) volumetric grids with different centers and scales. Namely, M-SDF of a single shape is a matrix X of dimension $n \times d$, where each row represent a single grid, and it can be fitted to a given shape’s SDF in less than 2 minutes with a single Nvidia A100 GPU. Furthermore, as X is a matrix representing a set it is compatible with existing expressive architectures, similar to the ones trained on points clouds [33]. Lastly, since the grids are centered near the shape’s surface M-SDF provides a parameter efficient representation.

We have used the M-SDF representation to train a Flow Matching model [21] on two main datasets: ShapeNetCoreV2 (3D Warehouse) [4] consisting of $\sim 50K$ 3D models divided to 55 classes and a dataset of $\sim 600K$ shapes and a matching text description [27]. We find that M-SDF to be the method of choice for 3D geometry representation and facilitate simple, generalizable, high quality, forward based 3D generative models.

2. Related work

Our work focuses on designing a representation that is beneficial for 3D generation. Below we review relevant previous works categorized by the type of representation they used for 3D shapes.

Grids and Triplanes. Several works suggested to train a diffusion model to predict a grid-based representation, which can be either 3D voxel grids [13, 19, 31], or using the Triplanes representation [10, 41, 45]. DiffRF [31] represents each shape as a volumetric radiance field on a voxel grid. [19] represents voxelized truncated signed distance functions (TSDF) encoded into a lower resolution features grid. Neural Wavelet [13] also advocates voxel grid structure containing wavelet coefficients of a truncated SDF signal. The main drawback of voxel grids is that they are restricted to a fixed and usually low resolution grids, mainly due to their cubic scaling.

The cubic scaling of volumetric grids motivated the Triplane representation [3, 39] using lower dimensional grids (1D and 2D) and encode a 3D function as a small MLP applied to sums of outer products. [10, 41, 45] utilize 2D diffusion model backbones on 3D training scenes represented by 2D feature planes. Despite the elegant correspondence to 2D images, preprocessing a large data set of 3D shapes into Triplane representation is compute intensive, requires to learn a shared MLP and in some cases a shared auto-encoder [10, 45]. Compared to our representation, grid-based representations contains many redundant empty voxels, since the shape’s surface usually occupied only a small fraction of the grid.

Neural Fields. Following the considerable success of Implicit Neural Representations (INRs) for modeling 3D shapes [29, 30, 34], several works suggest a generative model that produces a parameterization of an implicit function. [6, 51] employed a GAN to generate latent vector or volume representing an implicit function with a shared MLP. Shape-E [15] suggests a diffusion model that directly predicts the weights of a Multi-Layer Perceptron (MLP). However, training an MLP for each shape in a large dataset is compute intensive, and should consider the symmetries of MLP weights [32].

Point clouds. [26, 33, 48, 52] suggest to generate 3D Point clouds, taking advantage of existing permutation-equivariant architectures [36, 43]. Although point-clouds are easy to compute during preprocess and hence suitable for training on large datasets, post-processing it into a smooth surface results in loss of details or requires extensive number of points, which are hard to generate with permutation-equivariant models [44]. Point-E [33] suggests to train an additional upsampler diffusion model to scale the size of the generated point cloud, however it still fails to describe thin structures.

Other representations. There are papers that suggested methods that are related to our work and not fall into one of the categories above. [2] offers to generate a set of Volumetric Primitives [24] composed on top a coarse mesh, supervised with another 3D generative model’s outputs. [49] and [50] presented different representation for encoding occupancy field using a set of either structured or unstructured latent vectors. Similarly to our representation, they are compatible with the transformer architecture [43], trained as the generative model. However, all these methods require a significant compute intensive preprocessing stage for fitting a generalized representation for a large scale dataset. In contrast, our representation is both compact, has a simple structure, interpretable, and can be formed independently and quickly for each shape in the dataset.

3. Method

In this section we present our 3D shape representation based on an efficient approximation of the Signed Distance Function (SDF), we detail how it is computed for a dataset of shapes in a preprocess stage, and how it is used for training a flow-based generative model.

3.1. Mosaic-SDF Shape Representation

We advocate a simple and effective representation for the Signed Distance Functions (SDFs) of 3D shapes, suitable for generative models, that we call *Mosaic-SDF* (M-SDF). Our main focus is on building a representation that satisfies properties (i)-(iii) from Section 1.

Signed Distance Function (SDF). Our goal is to build a simple low parameter count approximation to the SDF of a shape $\mathcal{S} \subset \mathbb{R}^3$, where $\mathbf{x} \in \mathcal{S}$ are points *inside* the shape, $\mathbf{x} \notin \mathcal{S}$ are *outside* and $\mathbf{x} \in \partial\mathcal{S}$ are on the shape’s *boundary*. The SDF of the shape is defined by

$$F_{\mathcal{S}}(\mathbf{x}) = \begin{cases} -d(\mathbf{x}, \partial\mathcal{S}) & \mathbf{x} \in \mathcal{S} \\ d(\mathbf{x}, \partial\mathcal{S}) & \mathbf{x} \notin \mathcal{S} \end{cases}, \quad (1)$$

where the unsigned distance function is defined by $d(\mathbf{x}, \partial\mathcal{S}) = \min_{\mathbf{y} \in \partial\mathcal{S}} \|\mathbf{x} - \mathbf{y}\|_2$.

M-SDF representation. We approximate $F_{\mathcal{S}}$ with a volumetric function $F_X : \mathbb{R}^3 \rightarrow \mathbb{R}$ defined by the parameters X . The core idea is to define F_X as a weighted combination of a *set of local grids*. Following the three principles outlined above, (i)-(iii) we define the representation X to be the set of tuples

$$X = \{(\mathbf{p}_i, s_i, \mathbf{V}_i)\}_{i \in [n]}, \quad (2)$$




Figure 2. Mosaic-SDF representation.

where $[n] = \{1, \dots, n\}$, $\mathbf{p}_i \in \mathbb{R}^3$ are 3D point locations, $s_i \in \mathbb{R}$ are local scales, and $\mathbf{V}_i \in \mathbb{R}^{k \times k \times k}$ are local volumetric grids. See Figure 2 for an illustration of the local grids. We denote by $I_{\mathbf{V}_i} : [-1, 1]^3 \rightarrow \mathbb{R}$ the trilinear interpolants of the values \mathbf{V}_i over the origin-centered volumetric grid \mathbb{G} of the cube $[-1, 1]^3$ defined by

$$\mathbb{G} = \left\{ \frac{2(i_1, i_2, i_3) - n - 1}{n - 1} \mid i_1, i_2, i_3 \in [k] \right\}. \quad (3)$$

By convention the interpolants $I_{\mathbf{V}_i}$ vanish outside the cube $[-1, 1]^3$, *i.e.*, $I_{\mathbf{V}_i}(\mathbf{x}) = 0$ for $\mathbf{x} \notin [-1, 1]^3$. The parametric SDF approximation is then defined by

$$F_X(\mathbf{x}) = \sum_{i \in [n]} w_i(\mathbf{x}) I_{\mathbf{V}_i} \left(\frac{\mathbf{x} - \mathbf{p}_i}{s_i} \right) \quad (4)$$

where $w_i(\cdot)$ are scalar weight functions that define the contribution of the i ’th local grid; the $w_i(\cdot)$ are supported in the cube $[-1, 1]^3$ and satisfying partition of unity, *i.e.*,

$$\sum_{i \in [n]} w_i(\mathbf{x}) = 1, \quad \forall \mathbf{x} \in \mathbb{R}^3. \quad (5)$$

We opt for

$$w_i(\mathbf{x}) = \frac{\bar{w}_i(\mathbf{x})}{\sum_{j \in [n]} \bar{w}_j(\mathbf{x})}, \quad \bar{w}_i(\mathbf{x}) = \text{ReLU} \left[1 - \left\| \frac{\mathbf{x} - \mathbf{p}_i}{s_i} \right\|_{\infty} \right].$$

The domain of definition of F_X is the union of local grids,

$$\mathcal{D}(X) = \cup_{i \in [n]} \mathbb{B}_{\infty}(\mathbf{p}_i, s_i), \quad (6)$$

where $\mathbb{B}_{\infty}(\mathbf{p}, s) = \{\mathbf{x} \in \mathbb{R}^3 \mid \|\mathbf{x} - \mathbf{p}\|_{\infty} < s\}$ is the infinity ball of radius s centered at \mathbf{p} .

3.2. Computing M-SDF for Shape Dataset

Given an input shape \mathcal{S} (as defined above) from some dataset of shapes, we would like to compute its Mosaic-SDF representation X such that $F_X \approx F_{\mathcal{S}}$ in a neighborhood of the shape’s boundary surface, $\partial\mathcal{S}$. The computation of X consists of two stages: (i) *Initialization*: where we find X such that the domain of definition of F_X covers the surface of the shape, *i.e.*, $\partial\mathcal{S} \subset \mathcal{D}(X)$; and (ii) *Fine-tuning*: where we optimize X to improve approximation $F_X \approx F_{\mathcal{S}}$. This algorithm can be applied to each shape individually making it simple to parallelize and is computationally efficient compared to alternatives. The algorithm for converting a shape \mathcal{S} to its M-SDF representation is summarized in Algorithm 1. Figure 3 compares M-SDF representation and some of the popular existing representations for a *fixed budget* of parameters. Note that M-SDF provides the highest quality approximation while is only $\times 2$ slower than the fastest method, *i.e.*, the 3D volumetric grid. Later, in Section 4.2 we provide a more detailed evaluations and comparisons.




Figure 3. M-SDF representation: we compare the ground truth shape (left), with zero levelsets of (from left to right) M-SDF, Implicit Neural Representation (INR), Volumetric 3D Grid, and Triplane. All representations adhere the same budget of 355K parameters. Note that M-SDF provides the highest fidelity with an efficient computation time.

Initialization. We assume all shapes’ boundaries $\partial\mathcal{S}$ are provided in a way that allows basic operations like sampling and SDF computation, *e.g.*, as triangular meshes. We normalize all shapes so the $\partial\mathcal{S}$ fit in the cube $[-1, 1]^3$. We initialize the volumes centers $\{\mathbf{p}_i\}_{i \in [n]}$ using farthest point sampling [8] over the shape’s boundary $\partial\mathcal{S}$. Second, we set $s_i = s$ for all $i \in [n]$, where s is the minimal value that achieves a full coverage of the shape’s boundary, *i.e.*,

$$s = \min \{s > 0 \mid \partial\mathcal{S} \subset \mathcal{D}(X)\}, \quad (7)$$

where $\mathcal{D}(X)$ is defined in equation 6. To initialize the local volumes \mathbf{V}_i we simply store the corresponding values of the SDF F_S at the local grid coordinates, *i.e.*, $\mathbf{V}_i = F_S(\mathbf{p}_i + s_i \mathbb{G})$, where \mathbb{G} defined in equation 3.

Fine-Tuning. Although the initialization already provides a valid approximation to the shape’s SDF it can be further improved with quick fine-tuning. We optimize the initialized Mosaic-SDF representation X to reduce the following loss [9, 34] striving to regress the SDF’s values and first derivatives:

$$\mathcal{L}(X) = \mathcal{L}_1(X) + \lambda \mathcal{L}_2(X), \quad (8)$$

where

$$\mathcal{L}_1(X) = \frac{1}{m} \sum_{j \in [m]} \|F_X(\mathbf{x}_j) - F_S(\mathbf{x}_j)\|_1, \quad (9)$$

$$\mathcal{L}_2(X) = \frac{1}{m} \sum_{j \in [m]} \|\nabla_{\mathbf{x}} F_X(\mathbf{y}_j) - \nabla_{\mathbf{x}} F_S(\mathbf{y}_j)\|_2, \quad (10)$$

where $\|\cdot\|_1$, $\|\cdot\|_2$ represent the 1 and 2-norms (resp.), $\lambda > 0$ is a hyper-parameter. The sampling points $\{\mathbf{x}_i\}_{i \in [m]}$,

$\{\mathbf{y}_j\}_{j \in [m]}$ are spread uniformly over the shapes’ boundaries $\partial\mathcal{S}$ and their neighborhood; more details are in Section 4.1.

Algorithm 1: Mosaic-SDF preprocess.

Input : Shape \mathcal{S} , set size n , grid resolution k ,
 $\lambda \geq 0$
▷ Initialization
 $\{\mathbf{p}_i\}_{i \in [n]} \leftarrow$ farthest point sampling of $\partial\mathcal{S}$
 $\{s_i\}_{i \in [n]} \leftarrow s$ minimal covering scalar ▷ eq. 7
for $i \leftarrow 1$ **to** n **do**
 $\mathbf{V}_i \leftarrow F_S(\mathbf{p}_i + \mathbb{G} \cdot s_i)$ ▷ \mathbb{G} in eq. 3
▷ Fine-tuning
 $X \leftarrow \{(\mathbf{p}_i, s_i, \mathbf{V}_i)\}_{i \in [n]}$
while not converged do
 Take gradient step with $\nabla_X \mathcal{L}(X)$ ▷ eq. 8
Output: X

3.3. Mosaic-SDF Generation with Flow Matching

At this point we assume to be a given a dataset of N shapes paired with condition vectors (*e.g.*, classes or text embeddings), $\{(\mathcal{S}^i, \mathbf{c}^i)\}_{i \in [N]}$, and in a preprocess step we converted, using Algorithm 1, all shapes to M-SDF form, $\{(X^i, \mathbf{c}^i)\}_{i \in [N]}$. Our goal is to train a flow based generative model taking random noise to M-SDF samples. Given an M-SDF sample X the shape’s boundary can be extracted via zero level set contouring of F_X , *e.g.*, with Marching Cubes [25]. Below we recap the fundamentals of flow-based models adapted to our case and present the full training (Algorithm 2) and sampling (Algorithm 3).

Flow-based generative model. Flow-based generative models are designed to find a transformation taking samples from some simple noise distributions $X_0 \sim p(X_0)$ to conditioned data samples $X_1 \sim q(X_1|\mathbf{c})$. Our noise distribution is, as customary, a standard Gaussian $p = \mathcal{N}(0, I)$, and our empirical target distribution is

$$q(\cdot, \mathbf{c}) = \frac{1}{N} \sum_{i \in [N] | \mathbf{c}^i = \mathbf{c}} \delta(\cdot - X^i), \quad (11)$$

where $\delta(X)$ is a near-delta function (e.g., a Gaussian with small standard deviation σ) centered at 0. In our case, the data points are matrices,

$$X \in \mathbb{R}^{n \times d}, \text{ where } d = 3 + 1 + k^3, \quad (12)$$

with their row order considered only up to a permutation. A flow model is modeled with a *learnable velocity field* [5] denoted $U^\theta : [0, 1] \times \mathbb{R}^{n \times d} \times \mathbb{R}^c \rightarrow \mathbb{R}^{n \times d}$, where θ represents its learnable parameters, and c the dimension of an optional conditioning variable. Sampling from a flow model represented by U^θ is done by first randomizing $X_0 \sim p(X_0)$, and second, solving the ODE

$$\dot{X}_t = U_t^\theta(X_t, \mathbf{c}), \quad (13)$$

with initial condition X_0 , from time $t = 0$ until time $t = 1$, and $\mathbf{c} \in \mathbb{R}^c$ is the condition vector. Lastly, the desired sample of the model is defined to be X_1 .

Symmetric Data. As mentioned above $X \in \mathbb{R}^{n \times d}$ represents a set of n elements in \mathbb{R}^d (i.e., the elements are the rows of X). Differently put, permutation of the rows of X , i.e., $\mathbf{P}X$ with \mathbf{P} being a permutation matrix, is a *symmetry* of this representation, namely represents the same object $X \cong \mathbf{P}X$. Consequently, we would like our generative model to generate X and $\mathbf{P}X$ with the *same probability*. One way to achieve this is to consider noise density p that is *permutation invariant*, i.e.,

$$p(\mathbf{P}X) = p(X), \quad \text{for all } X, \mathbf{P}, \quad (14)$$

and a *permutation equivariant* flow field U^θ , i.e.,

$$U_t^\theta(\mathbf{P}X, \mathbf{c}) = \mathbf{P}U_t^\theta(X, \mathbf{c}), \quad \text{for all } t, X, \mathbf{c}, \mathbf{P}. \quad (15)$$

Indeed as proved in [18] (Theorem 1 and 2) equations 14 and 15 imply that the generations $X(1)$ using an equivariant model U^θ and invariant noise p are also permutation equivariant and $X(1)$ and $\mathbf{P}X(1)$ are generated with the same probability, as required. One benefit of this set symmetry is the existence of a powerful permutation equivariant neural architecture, namely a Transformer [33, 42] without positional encoding.

Algorithm 2: Flow Matching training.

Input : M-SDF dataset $\{X^i\}_{i \in [N]}$, p_{uncond} , σ
Initialize U_t^θ
while *not converged* **do**
 $t \sim \mathcal{U}([0, 1])$ ▷ sample time
 $(x_1, \mathbf{c}) \sim q(x_1, \mathbf{c})$ ▷ sample data and condition
 $\mathbf{c} \leftarrow \emptyset$ with probability p_{uncond} ▷ null condition
 $X_0 \sim p(X_0)$ ▷ sample noise
 $X_t \leftarrow tX_1 + (1 - \rho t)X_0$ ▷ eq. 17
 $\dot{X}_t \leftarrow X_1 - \rho X_0$
 Take gradient step on $\nabla_{\theta} \|U_t^\theta(X_t, \mathbf{c}) - \dot{X}_t\|^2$
Output: U^θ

Flow Matching. We use the recent formulation of Flow Matching (FM) [21] with its permutation equivariant variant [17]. Flow Matching models are similar to diffusion models in taking noise to data but directly regress the velocity field of a noise-to-data target flow and consequently have several benefits such as flexibility of noise-to-data paths, they are well defined for the entire time interval from noise to data (i.e., work with 0 SNR noise), easier to sample due to lower kinetic energy [40], and provide a competitive generation quality. We train Flow Matching with Classifier Free Guidance (CFG) [11] by minimizing the loss

$$\mathcal{L}(\theta) = \mathbb{E}_{t, b, p(X_0), q(X_1, \mathbf{c})} \left\| U_t^\theta(X_t, \bar{\mathbf{c}}(b)) - \dot{X}_t \right\|^2 \quad (16)$$

where $t \sim \mathcal{U}([0, 1])$ is the uniform distribution, $b \in \mathcal{B}(p_{\text{uncond}})$ is a Bernoulli random variable taking values 0, 1 with probability $(1 - p_{\text{uncond}}), p_{\text{uncond}}$ (resp.), $\bar{\mathbf{c}} = (1 - b) \cdot \mathbf{c} + b \cdot \emptyset$ where \emptyset is a symbol of null conditioning, and X_t is a path interpolating noise X_0 and data X_1 . We opt for paths X_t that form Optimal Transport displacement map [28] conditioned on a training sample $X_1 \sim q(X_1)$, i.e.,

$$X_t = (1 - \rho t)X_0 + tX_1, \quad \rho = 1 - \sigma \quad (17)$$

where $\sigma > 0$ is a hyper-parameter chosen to be $\sigma = 10^{-5}$ in our case. This path choice is referred to as Conditional Optimal Transport (Cond-OT) and it takes samples from $p(X_0)$ to samples from $\mathcal{N}(X_1, \sigma^2 I)$. Equivalent formulations of Flow Matching are also introduced in [1, 22].

4. Experiments

4.1. Implementation details

Datasets and preprocess. We train our generative model on two main datasets: 3D Warehouse data, commonly referred to as ShapenetCore v2 [4] consisting of 50K 3D polygonal meshes classified to 55 different categories, and a dataset of 600K polygonal meshes with matching text




Figure 4. Conditional samples from a Flow Matching model trained on M-SDF representations of ShapeNetCore v2.

Algorithm 3: ODE sampling.

Input : trained model U^θ , condition c ,
guidance parameter ω , number of steps m
 $X_0 \sim p(X_0)$ ▷ sample noise
 $h \leftarrow \frac{1}{m}$ ▷ step size
 $V_t(\cdot) \leftarrow (1 + \omega)U_t^\theta(\cdot|c) - \omega U_t^\theta(\cdot|\emptyset)$ ▷ CFG velocity
for $t = 0, h, \dots, 1 - h$ **do**
 $X_{t+h} \leftarrow \text{ODEStep}(V_t, X_t)$ ▷ ODE solver step
Output: X_1

descriptions [27]. We preprocess each polygonal mesh as follows: First, scale it to be bounded inside the unit cube $[-1, 1]^3$. Second, make it watertight using [12, 47]. Third, we run Algorithm 1 for 1K iterations, with $n = 1024$, $k = 7$, and $\lambda = 0.1$. Consequently the M-SDF tensor representation is $X \in \mathbb{R}^{1024 \times (3+1+7^3)}$ (see equation 12) and has 355K parameters in total. This last step takes less than 2 minutes on a single Nvidia A100 GPU (it takes a bit longer than the experiment in Figure 3 since here we used $\lambda > 0$). Fourth, noting that the M-SDF representation has three channels, (p_i, s_i, V_i) , we normalize p_i, s_i to have zero mean and unit max norm (using 50K random samples of each channel).

Flow Matching model architecture and training. We train a Flow Matching generative model [21] where for U^θ we use the transformer-based architecture [42] without positional encoding to achieve permutation equivariance, compatible with our M-SDF tensorial representation. Each element in the set (*i.e.*, row) of the noisy sample $X_t \in \mathbb{R}^{n \times d}$

is fed in as token, as well as the time t and the conditioning c . Our transformer is built with 24 layers with 16 heads and 1024 hidden dimension, which result in a 328M parameter model. We train U^θ for 500K iterations with batch size of 1024 using the ADAM optimizer [16] and learning rate of $1e-4$ with initial warm-up of 5K iterations. We additionally perform EMA (Exponential Moving Average) to the transformer’s weights. Both training were done on 8 nodes of 8 NVIDIA A100 GPUs, which takes around a week.

4.2. Representation evaluation

We start with comparing M-SDF to existing popular SDF representations used in 3D generative models focusing on two key aspects: preprocess efficiency, and parameter efficiency. We only consider SDF representations computed independently for each individual shape, *i.e.*, does not use latent space representations defined by a global encoder/decoder. The main reason for this choice is that all methods, including M-SDF, can be adapted to work on latent space, which is an orthogonal design choice. We compare to: 3D Volumetric Grid (3D-Grid), Triplane and Implicit Neural Representation (INR).

We consider 100 random (GT) models from ShapeNetCore-V2 (3D Warehouse) and for each SDF representation (M-SDF, 3D-Grid, Triplane, INR) we log its average preprocess time and surface approximation quality for varying parameter budget. The preprocess time is measured as the wall-clock time it takes a single Nvidia A100 GPU to compute the representation. For Triplane, INR and M-SDF we use the loss in equation 8 with $\lambda = 0$; 3D-Grid is computed by evaluating the GT SDF at the grid nodes. The surface approximation quality

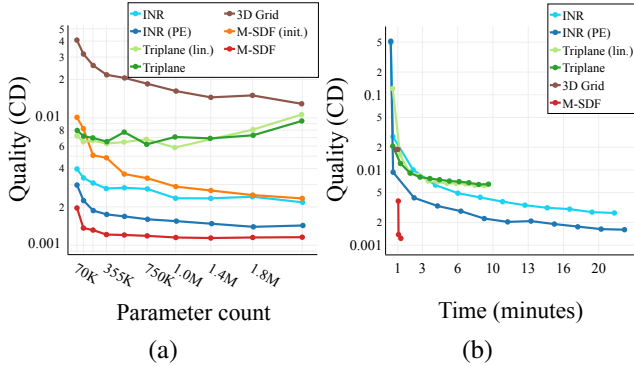


Figure 5. (a) 3D approximation quality vs. representation parameter count; (b) pre-process training time vs. 3D approximation quality for a fixed representation budget of 355K parameters.

is measured by the Chamfer Distance (CD) between the extracted surface mesh from the SDF representation and the GT mesh. Figure 5 summarizes the results in two graphs: (a) shows the surface approximation quality of each representation for different parameter budgets; and (b) surface approximation quality versus preprocess time of each representation. For INR we additionally examine the option of using Positional Encoding (PE) [30], which incorporates high frequencies as input to the MLP. For Triplane we evaluated the two alternatives of aggregating the projected features, as suggested in [39], using learned linear decoder (denoted as lin.) or using a small MLP decoder. We note that while previous works that use the Triplane representation [3, 41] suggested additional regularizations, we tested Triplane with the same supervision and loss as the other methods that require optimization, namely INR and M-SDF. This can potentially explain the degradation in the approximation quality as the parameter count increases. Additionally, for M-SDF we report the surface approximation quality both at initialization and after fine-tuning (see Section 3.2). As can be seen in the graphs, M-SDF is superior to INR in terms of surface approximation per parameter budget while is computable in only a fraction of the time. 3D Grids are the only faster baseline to M-SDF but their approximation-parameter trade-off is considerably worse (see also Figure 3).

4.3. Class conditional generation

In this section we evaluate our class-conditioned generative FM model trained on the ShapeNetCore-V2 (3D Warehouse) [4] where the 55 classes of objects in this dataset are used as conditions. We follow the split to train/val/test suggested by [49]. Following and extending protocols of past works, Table 1 reports quantitative evaluation for the 5 largest classes (containing over 3K shapes each) comparing M-SDF and relevant recent baselines.

| | FPD (\downarrow) | KPD (\downarrow) | COV (\uparrow ,%) | | MMD (\downarrow) | | 1-NNA (\downarrow ,%) | |
|--------------|----------------------|----------------------|----------------------|-------|----------------------|-------|--------------------------|-------|
| | | | CD | EMD | CD | EMD | CD | EMD |
| airplane | | | | | | | | |
| 3DILG | 1.83 | 3.22 | 41.09 | 32.67 | 4.69 | 4.73 | 82.67 | 84.41 |
| NW | 0.81 | 1.26 | 51.98 | 45.05 | 3.36 | 4.19 | 68.32 | 73.76 |
| S2VS | 0.94 | 1.65 | 51.98 | 40.59 | 3.80 | 4.45 | 69.06 | 76.73 |
| Ours | 0.44 | 0.50 | 52.48 | 51.49 | 3.54 | 3.78 | 62.62 | 69.55 |
| car | | | | | | | | |
| 3DILG | 2.84 | 6.24 | 18.86 | 20.57 | 4.67 | 3.83 | 93.43 | 90.57 |
| NW | - | - | - | - | - | - | - | - |
| S2VS | 1.32 | 2.17 | 37.71 | 40.00 | 4.13 | 3.52 | 84.57 | 86.57 |
| Ours | 0.46 | 0.48 | 45.71 | 51.43 | 2.87 | 2.75 | 70.00 | 66.00 |
| chair | | | | | | | | |
| 3DILG | 1.64 | 2.00 | 37.87 | 39.94 | 20.37 | 10.54 | 74.11 | 69.38 |
| NW | 1.41 | 1.29 | 43.79 | 47.04 | 16.53 | 9.47 | 59.47 | 64.20 |
| S2VS | 0.77 | 0.63 | 51.78 | 52.37 | 16.97 | 9.44 | 58.43 | 60.80 |
| Ours | 0.52 | 0.19 | 48.22 | 55.03 | 15.47 | 9.13 | 51.04 | 55.62 |
| sofa | | | | | | | | |
| 3DILG | 3.19 | 5.83 | 25.95 | 29.11 | 26.41 | 10.71 | 84.81 | 77.85 |
| NW | - | - | - | - | - | - | - | - |
| S2VS | 1.17 | 1.70 | 48.73 | 51.90 | 10.83 | 7.25 | 62.66 | 57.91 |
| Ours | 0.63 | 0.62 | 46.20 | 48.10 | 12.43 | 7.60 | 61.71 | 55.70 |
| table | | | | | | | | |
| 3DILG | 2.86 | 4.13 | 29.45 | 30.88 | 22.96 | 10.18 | 78.27 | 78.74 |
| NW | 1.49 | 2.20 | 51.07 | 47.98 | 13.27 | 7.72 | 56.41 | 58.67 |
| S2VS | 0.83 | 0.92 | 53.44 | 49.41 | 14.06 | 8.01 | 59.74 | 61.05 |
| Ours | 0.47 | 0.21 | 52.97 | 53.21 | 13.49 | 7.74 | 51.31 | 50.59 |

Table 1. Evaluation of our class conditioning generation model trained on 3D Warehouse [14] compared to baselines. We report results on the 5 largest classes in the dataset. KPD is multiplied by 10^3 , MMD-CD by 10^3 and MMD-EMD by 10^2 .




Figure 6. Class conditioning generation of 3D shapes compared to relevant baselines. Note the high fidelity demonstrated in M-SDF generated shapes compared to the (overly smooth) baselines.

Evaluation metrics. We make use of several standard metrics for evaluating the performance of our M-SDF generative models. All these metrics quantify differences between a set of reference shapes S_r and generated shapes S_g . To measure distances between shapes we follow previous works (e.g., [26, 48, 50]) and use the Chamfer Distance (CD) and Earth Moving Distance (EMD). Using these distances we compute: Maximum Mean Discrepancy (MMD), Coverage (COV), and 1-nearest neighbor accuracy (1-NNA) to quantify fidelity, diversity and distributional similarity, respectively. Furthermore, following [33, 50] we use the 3D analogs of the Fréchet Inception Distance (FID)




Figure 7. Text-to-3D samples from a Flow Matching model trained on M-SDF representations.

and Kernel Inception Distance (KID), commonly used to evaluate image generative models. We employ a pre-trained PointNet++ [37] to extract features. As in [50] we will refer to these metrics as Fréchet PointNet++ Distance (FPD) and Kernel PointNet++ Distance (KPD). Additional information and implementation details are provided in the supplementary material.

Baselines. We compare to Neural Wavelet (NW) [13] that generate a grid-based representation. Note that this works trains an unconditional model for each class separately, making it arguably simpler than a single conditioned model on the entire 55 classes. We additionally compare to 3DILG [49] and 3DShape2VecSet (S2VS) [50], that suggest to generate structured or unstructured latent vectors, encoding an occupancy field using a transformer.

As can be read in Table 1 our M-SDF based generative model compares favorably to the baselines, achieving best, or second best performance in all metrics. Figure 6 shows qualitative comparison for two classes common for all baselines, *i.e.*, chairs and airplanes. Note that generation with M-SDF allows higher fidelity and sharper surfaces compared to baselines which tend to be overly smooth. Figure 4 shows additional class conditional generations with our M-SDF model.

SDF evaluation time.

An additional advantage of M-SDF compared to other performant baselines is the relative efficiency and flexibility in extracting the zero level set of the SDF F_X (equation 4). 3DILG and S2VS require forward pass in a transformer for function evaluation. NW is restricted to generate a 3D grid in a single resolution. M-SDF can be evaluated efficiently by querying only the relevant local grids. In Table 2 we log, for each method, the total time of extracting the zero levelset of the SDF (with Marching Cubes) using cubic grids

| | 64^3 | 128^3 | 256^3 | 512^3 |
|--------------|--------|---------|---------|---------|
| 3DILG | 0.3 | 2.33 | 18.44 | 159.56 |
| NW | - | - | 0.61 | - |
| S2VS | 0.06 | 0.36 | OOM | OOM |
| Ours | 0.05 | 0.34 | 2.74 | 21.48 |

Table 2. Surface extraction time (in seconds).

of different resolutions (64^3 , 128^3 , 256^3 , and 512^3); OOM stands for Out of Memory.

4.4. Text-to-3D generation

Lastly we provide a qualitative evaluation of our M-SDF based generative FM model trained on a dataset of 600K shapes with matching text captions [27]. We utilize a pre-trained text model [38] as our textual embedding, passing this embedding to the model as conditioning. Figure 7 depicts pairs of generated shapes and the text conditions used to generate it.

5. Summary and Future Work

We presented a novel 3D shape representation, Mosaic-SDF, that is geared towards 3D generative models and offers a simple and efficient preprocessing, favorable approximation-parameter trade-off, and a simple tensorial structure compatible with powerful modern architectures (*i.e.*, transformers). We have used M-SDF to train Flow Matching generative models and demonstrated state of the art results for forward-based models. We believe that M-SDF is the method of choice for 3D generation however still possesses some limitations and can be extended in several ways: First, currently we only encode the SDF, missing texture/color/light information. An interesting extension is to incorporate texture and/or light field data. Second, in our architecture we use a simple linear layer passing the local grids into the transformer. A possible extension here is to incorporate convolution layers and/or autoencoders to further increase resolution/data reuse of the representation. Lastly, making M-SDF equivariant to orientations, *e.g.*, by adding local coordinate frames, can improve the generalization abilities of the trained model, which is currently only permutation equivariant.

Acknowledgements

We thank Itai Gat, Matt Le, Ricky Chen for valuable advice and discussion, and Biao Zhang for sharing evaluation code. OP is supported by a grant from Israel CHE Program for Data Science Research Centers.

References

- [1] Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022. 5
- [2] Mallikarjun B R, Xingang Pan, Mohamed Elgharib, and Christian Theobalt. Gvp: Generative volumetric primitives, 2023. 3
- [3] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *arXiv*, 2021. 2, 7
- [4] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015. 2, 5, 7
- [5] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018. 5
- [6] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [7] Wikipedia contributors. Janus — wikipedia, the free encyclopedia, 2010. [2023]. 1
- [8] Yuval Eldar, Michael Lindenbaum, Moshe Porat, and Yehoshua Y Zeevi. The farthest point strategy for progressive image sampling. *IEEE Transactions on Image Processing*, 6(9):1305–1315, 1997. 4
- [9] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *Proceedings of Machine Learning and Systems 2020*, pages 3569–3579. 2020. 4
- [10] Anchit Gupta, Wenhan Xiong, Yixin Nie, Ian Jones, and Barlas Oğuz. 3dgen: Triplane latent diffusion for textured mesh generation, 2023. 1, 2
- [11] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 5
- [12] Jingwei Huang, Hao Su, and Leonidas Guibas. Robust watertight manifold surface generation method for shapenet models. *arXiv preprint arXiv:1802.01698*, 2018. 6
- [13] Ka-Hei Hui, Ruihui Li, Jingyu Hu, and Chi-Wing Fu. Neural wavelet-domain diffusion for 3d shape generation. 2022. 2, 8
- [14] Trimble Inc. 3d warehouse. 7
- [15] Heewoo Jun and Alex Nichol. Shap-e: Generating conditional 3d implicit functions, 2023. 1, 2
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [17] Leon Klein, Andreas Krämer, and Frank Noé. Equivariant flow matching. *arXiv preprint arXiv:2306.15030*, 2023. 5
- [18] Jonas Köhler, Leon Klein, and Frank Noé. Equivariant flows: exact likelihood generative learning for symmetric densities. In *International conference on machine learning*, pages 5361–5370. PMLR, 2020. 5
- [19] Muheng Li, Yueqi Duan, Jie Zhou, and Jiwen Lu. Diffusion-sdf: Text-to-shape via voxelized diffusion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 2
- [20] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiao-hui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1
- [21] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023. 2, 5, 6
- [22] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022. 5
- [23] Zhen Liu, Yao Feng, Michael J. Black, Derek Nowrouzezahrai, Liam Paull, and Weiyang Liu. Meshdiffusion: Score-based generative 3d mesh modeling. In *International Conference on Learning Representations*, 2023. 1
- [24] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhoefer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering. *ACM Trans. Graph.*, 40(4), 2021. 3
- [25] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pages 347–353. 1998. 4
- [26] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 7
- [27] Tiange Luo, Chris Rockwell, Honglak Lee, and Justin Johnson. Scalable 3d captioning with pretrained models. *arXiv preprint arXiv:2306.07279*, 2023. 2, 6, 8
- [28] Robert J McCann. A convexity principle for interacting gases. *Advances in mathematics*, 128(1):153–179, 1997. 5
- [29] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [30] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 2, 7
- [31] Norman Müller, Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Buló, Peter Kotschieder, and Matthias Nießner. Diffri: Rendering-guided 3d radiance field diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4328–4338, 2023. 1, 2

- [32] Aviv Navon, Aviv Shamsian, Idan Achituve, Ethan Fetaya, Gal Chechik, and Haggai Maron. Equivariant architectures for learning in deep weight spaces. *arXiv preprint arXiv:2301.12780*, 2023. 2
- [33] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts, 2022. 1, 2, 5, 7
- [34] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 4
- [35] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv*, 2022. 1
- [36] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *arXiv preprint arXiv:1612.00593*, 2016. 2
- [37] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017. 8
- [38] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020. 8
- [39] Sara Fridovich-Keil and Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *CVPR*, 2023. 2, 7
- [40] Neta Shaul, Ricky TQ Chen, Maximilian Nickel, Matthew Le, and Yaron Lipman. On kinetic optimal probability paths for generative models. In *International Conference on Machine Learning*, pages 30883–30907. PMLR, 2023. 5
- [41] J. Ryan Shue, Eric Ryan Chan, Ryan Po, Zachary Ankner, Jiajun Wu, and Gordon Wetzstein. 3d neural field generation using triplane diffusion, 2022. 1, 2, 7
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017. 5, 6
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2017. 2, 3
- [44] Edward Wagstaff, Fabian B Fuchs, Martin Engelcke, Michael A Osborne, and Ingmar Posner. Universal approximation of functions on sets. *The Journal of Machine Learning Research*, 23(1):6762–6817, 2022. 2
- [45] Tengfei Wang, Bo Zhang, Ting Zhang, Shuyang Gu, Jianmin Bao, Tadas Baltrusaitis, Jingjing Shen, Dong Chen, Fang Wen, Qifeng Chen, and Baining Guo. Rodin: A generative model for sculpting 3d digital avatars using diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4563–4573, 2023. 1, 2
- [46] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *arXiv preprint arXiv:2305.16213*, 2023. 1
- [47] Francis Williams. Point cloud utils, 2022. <https://www.github.com/fwilliams/point-cloud-utils>. 6
- [48] Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. Lion: Latent point diffusion models for 3d shape generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 1, 2, 7
- [49] Biao Zhang, Matthias Nießner, and Peter Wonka. 3DILG: Irregular latent grids for 3d generative modeling. In *Advances in Neural Information Processing Systems*, 2022. 3, 7, 8
- [50] Biao Zhang, Jiapeng Tang, Matthias Niessner, and Peter Wonka. 3dshape2vecset: A 3d shape representation for neural fields and generative diffusion models, 2023. 3, 7, 8
- [51] Xin-Yang Zheng, Yang Liu, Peng-Shuai Wang, and Xin Tong. Sdf-stylegan: Implicit sdf-based stylegan for 3d shape generation. In *Comput. Graph. Forum (SGP)*, 2022. 2
- [52] Linqi Zhou, Yilun Du, and Jiajun Wu. 3d shape generation and completion through point-voxel diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5826–5835, 2021. 1, 2