

EDGERUNNER: AUTO-REGRESSIVE AUTO-ENCODER FOR ARTISTIC MESH GENERATION

Anonymous authors

Paper under double-blind review

<https://edgerunnerpaper.github.io/>

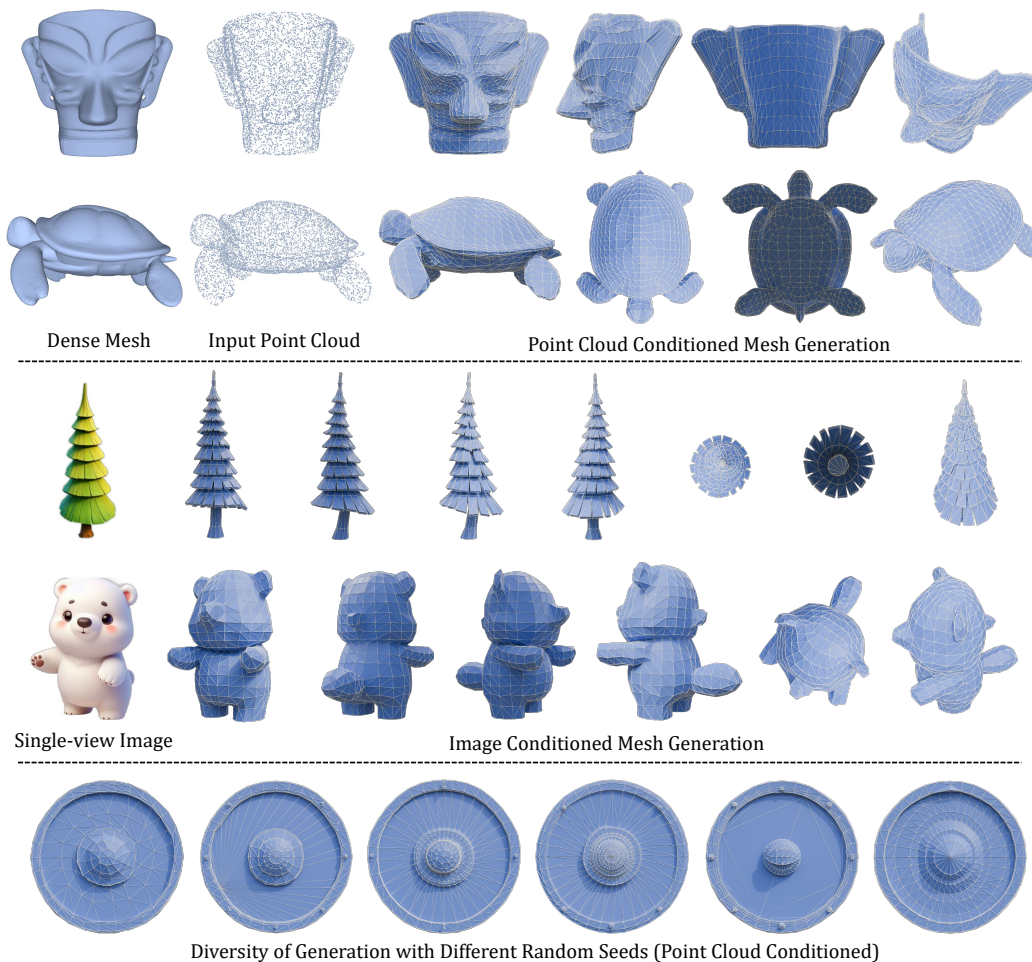


Figure 1: **EdgeRunner** efficiently generates diverse, high-quality artistic meshes conditioned on point clouds or single-view images.

ABSTRACT

Current auto-regressive mesh generation methods suffer from issues such as incompleteness, insufficient detail, and poor generalization. In this paper, we propose an Auto-regressive Auto-encoder (ArAE) model capable of generating high-quality 3D meshes with up to 4,000 faces at a spatial resolution of 512^3 . We introduce a novel mesh tokenization algorithm that efficiently compresses triangular meshes into 1D token sequences, significantly enhancing training efficiency. Furthermore, our model compresses variable-length triangular meshes into a fixed-length latent space, enabling training latent diffusion models for better generalization. Extensive experiments demonstrate the superior quality, diversity, and generalization capabilities of our model in both point cloud and image-conditioned mesh generation tasks.

1 INTRODUCTION

Automatic 3D content generation, particularly the generation of widely used polygonal meshes, holds the potential to revolutionize industries such as digital gaming, virtual reality, and filmmaking. Generative models can make 3D asset creation more accessible to non-experts by drastically reducing the time and effort involved. This democratization opens up opportunities for a wider range of individuals to contribute to and innovate within the 3D content space, fostering greater creativity and efficiency across these sectors.

Previous research on 3D generation has explored a variety of approaches. For example, optimization-based methods, such as using score distillation sampling (SDS) (Poole et al., 2022; Lin et al., 2023; Liu et al., 2023b; Tang et al., 2023a), lift 2D diffusion priors into 3D without requiring any 3D data. In contrast, large reconstruction models (LRM) (Hong et al., 2023; Wang et al., 2023b; Xu et al., 2023b; Li et al., 2023; Weng et al., 2024b) directly train feed-forward models to predict neural radiance fields (NeRF) or Gaussian Splatting from single or multi-view image inputs. Lastly, 3D-native latent diffusion models (Zhang et al., 2024c; Wu et al., 2024b; Li et al., 2024c) encode 3D assets into latent spaces and generate diverse contents by performing diffusion steps in the latent space. However, all these approaches rely on continuous 3D representations, such as NeRF or SDF grids, which lose the discrete face indices in triangular meshes during conversion. Consequently, they require post-processing, such as marching cubes (Lorensen & Cline, 1998) and re-meshing algorithms, to extract triangular meshes. These meshes differ significantly from artist-created ones, which are more concise, symmetric, and aesthetically structured. Additionally, these methods are limited to generating watertight meshes and cannot produce single-layered surfaces.

Recently, several approaches (Siddiqui et al., 2024a; Chen et al., 2024d;b; Weng et al., 2024a; Chen et al., 2024e) have attempted to tokenize meshes into 1D sequences and leverage auto-regressive models for direct mesh generation. Specifically, MeshGPT (Siddiqui et al., 2024a) proposes to empirically sort the triangular faces and apply a vector-quantization variational auto-encoder (VQ-VAE) to tokenize the mesh. MeshXL (Chen et al., 2024b) directly flattens the vertex coordinates and does not use any compression other than vertex discretization. Since these methods directly learn from mesh vertices and faces, they can preserve the topology information and generate artistic meshes. However, these auto-regressive mesh generation approaches still face several challenges. (1) Generation of a large number of faces: due to the inefficient face tokenization algorithms, most prior methods can only generate meshes with fewer than 1,600 faces, which is insufficient for representing complex objects. (2) Generation of high-resolution surface: previous works quantize mesh vertices to a discrete grid of only 128^3 resolution, which results in significant accuracy loss and unsmooth surfaces. (3) Model generalization: training auto-regressive models with difficult input modalities is challenging. Previous approaches often struggle to generalize beyond the training domain when conditioning on single-view images.

In this paper, we present a novel approach named **EdgeRunner** to address the aforementioned challenges. Firstly, we introduce a mesh tokenization method based on EdgeBreaker (Rossignac, 1999) that compresses sequence length by 50% and reduces long-range dependency between tokens, significantly improving the training efficiency. Secondly, we propose an Auto-regressive Auto-encoder (ArAE) that compresses variable-length triangular meshes into fixed-length latent codes. This latent space can be used to train latent diffusion models conditioned on other modalities, offering better generalization capabilities. We also enhance the training pipeline to support higher quantization resolution. These improvements enable EdgeRunner to generate diverse, high-quality artistic meshes with up to 4,000 faces and vertices discretized at a resolution of 512^3 — resulting in sequences that are twice as long and four times higher in resolution compared to previous methods.

In summary, our contributions are as follows:

1. We introduce a novel mesh tokenization algorithm, adapted from EdgeBreaker, which supports lossless face compression, prevents flipped faces, and reduces long-range dependencies to facilitate learning.
2. We propose an Auto-regressive Auto-encoder (ArAE), comprising a lightweight encoder and an auto-regressive decoder, capable of compressing variable-length triangular meshes into fixed-length latent codes.

3. We demonstrate that the latent space of ArAE can be leveraged to train latent diffusion models for better generalization, enabling conditioning of different input modalities such as single-view images.
4. Extensive experiments show that our method generates high-quality and diverse artistic meshes from point clouds or single-view images, exhibiting improved generalization and robustness compared to previous methods.

2 RELATED WORK

2.1 OPTIMIZATION-BASED 3D GENERATION

Early 3D generation methods relied on SDS-based optimization techniques (Jain et al., 2022; Poole et al., 2022; Wang et al., 2023a; Mohammad Khalid et al., 2022; Michel et al., 2022) due to limited 3D data. Subsequent works advanced in generation quality (Lin et al., 2023; Wang et al., 2023d; Chen et al., 2023c;e; Sun et al., 2023; Qiu et al., 2024), reducing generation time (Tang et al., 2023a; Yi et al., 2023; Lorraine et al., 2023; Xu et al., 2024a), enabling 3D editing (Zhuang et al., 2023; Singer et al., 2023; Raj et al., 2023; Chen et al., 2024c), and conditioning on images (Xu et al., 2023a; Tang et al., 2023b; Melas-Kyriazi et al., 2023; Liu et al., 2023b; Qian et al., 2023; Shi et al., 2023). Other approaches first predict multi-view images, then apply reconstruction algorithms to generate the final 3D models (Long et al., 2023; Li et al., 2024b; Pang et al., 2024; Tang et al., 2024b). Recently, Unique3D (Wu et al., 2024a) introduced a method that combines high-resolution multi-view diffusion models with an efficient mesh reconstruction algorithm, achieving both high quality and fast image-to-3D generation.

2.2 FEED-FORWARD 3D GENERATION

With the introduction of large-scale datasets (Deitke et al., 2023b;a), more recent works propose to use feed-forward 3D models. The Large Reconstruction Model (LRM)(Hong et al., 2023) demonstrated that end-to-end training of a triplane-NeRF regression model scales effectively to large datasets and generates 3D assets within seconds. While LRM significantly accelerates generation speed, the resulting meshes often exhibit lower quality and a lack of diversity. Subsequent research has sought to improve generation quality by incorporating multi-view images as inputs(Xu et al., 2024b; Li et al., 2023; Wang et al., 2023b; He & Wang, 2023; Siddiqui et al., 2024b; Xie et al., 2024; Wang et al., 2024) and by adopting more efficient 3D representations (Zhang et al., 2024a; Li et al., 2024a; Wei et al., 2024; Zou et al., 2023; Tang et al., 2024a; Xu et al., 2024d; Zhang et al., 2024b; Chen et al., 2024a; Yi et al., 2024).

2.3 DIFFUSION-BASED 3D GENERATION

Analogous to 2D diffusion models for image generation, significant efforts have been made to develop 3D-native diffusion models capable of conditional 3D generation. Early approaches typically rely on uncompressed 3D representations, such as point clouds, NeRFs, tetrahedral grids, and volumes (Nichol et al., 2022; Jun & Nichol, 2023; Gupta et al., 2023; Cheng et al., 2023; Ntavelis et al., 2023; Zheng et al., 2023; Zhang et al., 2023; Liu et al., 2023c; Müller et al., 2023; Chen et al., 2023d; Cao et al., 2023; Chen et al., 2023a; Wang et al., 2023c; Yariv et al., 2023; Liu et al., 2023a; Xu et al., 2024c; Yan et al., 2024) to train diffusion models. However, these methods are often limited by small-scale datasets and struggle with generalization or producing high-quality assets. More recent approaches have focused on adapting latent diffusion models to 3D (Zhao et al., 2023; Zhang et al., 2024c; Wu et al., 2024b; Li et al., 2024c; Lan et al., 2024; Hong et al., 2024; Tang et al., 2023c; Chen et al., 2024f). These methods first train a VAE to compress 3D representations into a more compact form, which enables more efficient diffusion model training. Unlike the straightforward image representations in 2D, 3D latent diffusion models involve numerous design choices, leading to varied performance outcomes. For example, CLAY (Zhang et al., 2024c) has demonstrated that a transformer-based 3D latent diffusion model can scale to large datasets and generalize well across diverse input conditions.

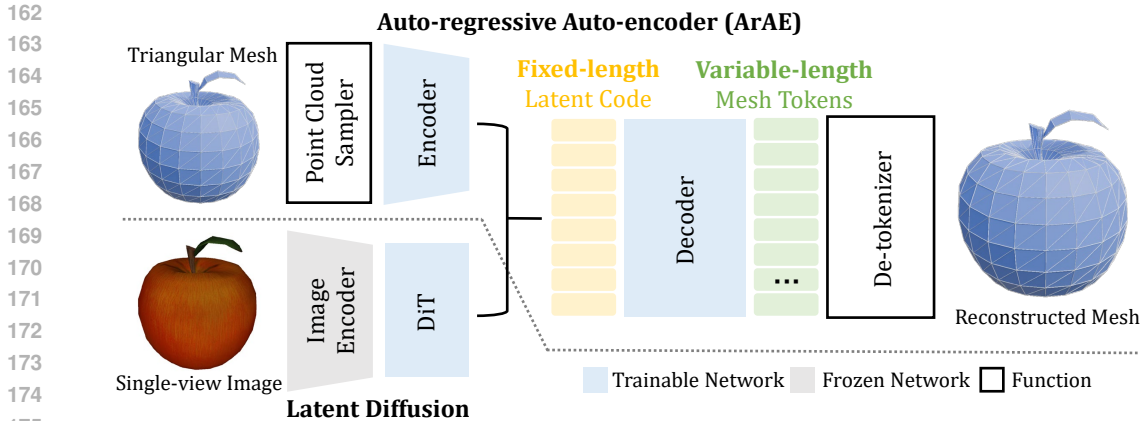


Figure 2: **Pipeline of our method.** Our ArAE model compresses variable-length mesh into fixed-length latent code, which can be further used to train latent diffusion models conditioned on other input modalities, such as single-view images.

2.4 AUTO-REGRESSIVE MESH GENERATION

The above works require additional post-processing steps to extract triangular meshes and fail to model the mesh topology. Recently, approaches using auto-regressive models to directly generate meshes have emerged. MeshGPT (Siddiqui et al., 2024a) pioneered this approach by tokenizing a mesh through face sorting and compression with a VQ-VAE, followed by using an auto-regressive transformer to predict the token sequence. This method allows for the generation of meshes with varying face counts and incorporates direct supervision from topology information, which is often overlooked in other approaches. Subsequent works (Chen et al., 2024b; Weng et al., 2024a; Chen et al., 2024d) have explored different model architectures and extended this approach to conditional generation tasks, such as point cloud generation. However, these methods are limited to meshes with fewer than 1,000 faces due to the computational cost of mesh tokenization and exhibit limited generalization capabilities. A concurrent work, MeshAnythingV2 (Chen et al., 2024e), introduces an improved mesh tokenization technique, increasing the maximum number of faces to 1,600. Our approach also falls under the category of auto-regressive mesh generation but aims to further extend the maximum face count and provide control over the target face number during inference.

3 EDGERUNNER

3.1 COMPACT MESH TOKENIZATION

Auto-regressive models process information in the form of discrete token sequences. Thus, compact tokenization is crucial as it allows information to be represented with fewer tokens accurately. For example, text tokenizers have been a central research direction for large language models (LLMs). The GPT and Llama series (Touvron et al., 2023; Brown et al., 2020) utilize the byte-pair encoding (BPE) tokenizer, which combines sub-word units into single tokens for highly compact and lossless compression.

In contrast, tokenization techniques used in prior auto-regressive mesh generation works mainly suffer from two issues. Some prior works use *lossy* VQ-VAEs (Siddiqui et al., 2024a; Chen et al., 2024d; Weng et al., 2024a), which sacrifices the mesh generation quality. Others opt for *zero-compression* by not using face tokenizers (Chen et al., 2024b), which poses training challenges due to the inefficiency.

In this paper, we introduce a tokenization scheme that allows us to represent a mesh compactly and efficiently, which is based on the well-established triangular mesh compression algorithm EdgeBreaker (Rossignac, 1999). **The key insight for mesh compression is to maximize edge sharing between adjacent triangles.** By sharing an edge with the previous triangle, the next triangle re-

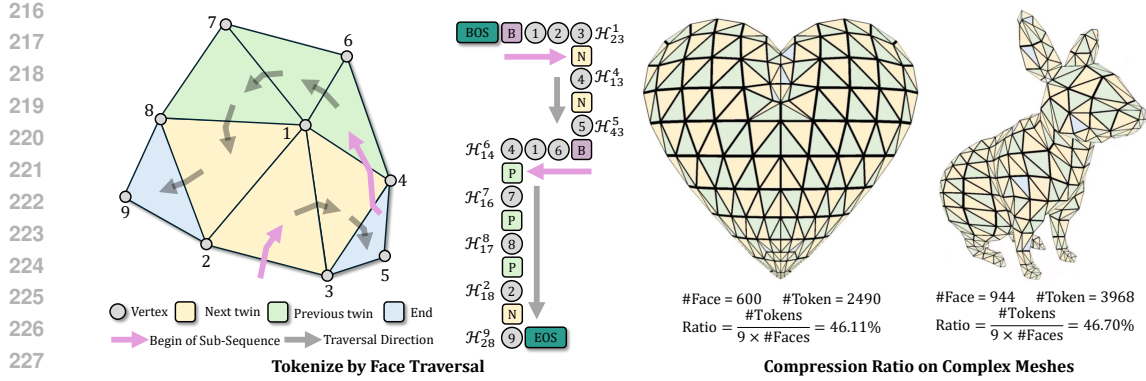


Figure 3: **Illustration of our mesh tokenizer.** Our tokenizer traverses the 3D mesh triangle-by-triangle and converts it into a 1D token sequence. Through edge sharing, we reach a compression rate of 50% (4 or 5 tokens per face on average) compared to naïve tokenization of 9 tokens per face.

quires only one additional vertex instead of three. We illustrate our mesh tokenization process with an example below, and provide more details in the appendix.

Half-edge. EdgeBreaker (Rossignac, 1999) uses half-edge data structures (Weiler, 1986) for triangular face traversal. An illustration is provided in Figure 4. We use \mathcal{H}_i to denote a half-edge. For example, \mathcal{H}_{41}^3 is the half-edge pointing from vertex 4 to 1, with vertex 3 across the face. Starting from \mathcal{H}_{41}^3 , we can traverse to the *next* half-edge \mathcal{H}_{13}^4 and the *next twin* half-edge \mathcal{H}_{31}^2 . Reversely, the *previous* half-edge is \mathcal{H}_{34}^1 and the *previous twin* half-edge is \mathcal{H}_{43}^5 . The half-edge data structure has also been used in recent learning-based mesh generation work (Shen et al., 2024).

Vertex Tokenization. To tokenize a mesh into a discrete sequence, vertex coordinates require discretization. Following previous works (Siddiqui et al., 2024a), we normalize the mesh to a unit cube and quantize the continuous vertex coordinates into integers according to a quantization resolution, which is 512 in this work. Each vertex is therefore represented by three integer coordinates, which are then flattened in XYZ order as tokens. With some abuse of notion, we denote the XYZ tokens as a single vertex token using \bullet .

Face Tokenization. We traverse through all faces following the half-edges. To illustrate the process, we use the mesh example in Figure 3. The process starts with one half-edge, where \mathcal{H}_{23}^1 is picked as the beginning of the current traversal. We signify the start of a traversal as **B**. We then append the vertex across the half-edge **1** as the first vertex token. Within the same triangular face, the two vertices **2** **3** are also appended following the direction of \mathcal{H}_{23}^1 .

During traversal, we visit the *next twin* half-edge whenever possible, and only reverse the half-edge direction to the *previous twin* half-edge when we exhaust all triangles in the current traversal. Returning to the example in Figure 3, we follow \mathcal{H}_{23}^1 and reach \mathcal{H}_{13}^4 . Thus, we append **N** to signify the *next twin* traversal direction and we only need to append **4** as **1** **3** are shared. The same process is repeated for \mathcal{H}_{43}^5 with **N** **5** added to the current sub-sequence.

We have completed the current traversal as no adjacent faces can be found for \mathcal{H}_{43}^5 . The sub-sequence for the current traversal is thus **B** **1** **2** **3** **N** **4** **N** **5**.

To begin a new sub-sequence, we reversely retrieve the last-added half-edges to traverse in the opposite directions. As the last-added half-edge \mathcal{H}_{43}^5 doesn't have any adjacent faces, we skip it and instead consider \mathcal{H}_{13}^4 . We go opposite to its *previous twin* half-edge \mathcal{H}_{14}^6 . As this is a new sub-sequence, **B** **6** **1** **4** are added.

We continue finding the un-visited faces in the neighborhood of \mathcal{H}_{14}^6 and arriving at its *previous twin* half-edge \mathcal{H}_{16}^7 . Thus, we add **P** **7** to the current sub-mesh sequence as **6** **1** are shared. The




Figure 4: Half-edge representation for triangular faces.

process is repeated and $\text{P } 8 \text{ P } 2 \text{ N } 9$ are added. As all triangular faces have been visited, the face tokenization process of the mesh is complete.

Auxiliary Tokens. Similar to prior work in LLMs, we prepend a BOS at the beginning of a mesh sequence and append a EOS at the end of a mesh sequence.

Detokenization. It is straightforward to reconstruct the original mesh from a mesh token sequence. We iterate over the tokens while maintaining a state machine. Each B of a sub-sequence is always followed by three vertex tokens. Each N or P is followed by a single vertex token, and we retrieve two previous vertex tokens based on the traversal direction to reconstruct the triangle. Finally, we merge duplicate vertices, as they may appear multiple times from different sub-sequences, and output the reconstructed mesh.

Advantages. Our tokenizer benefits model training in several ways: (1) Each face requires an average of 4 to 5 tokens, achieving approximately 50% compression compared to the 9 tokens used in previous works (Chen et al., 2024b;d). This increased efficiency enables the model to generate more faces with the same number of tokens and facilitates training on datasets containing a higher number of faces. (2) Our traversal is designed to avoid long-range dependency between tokens. Each token only relies on a short context of previous tokens, which further mitigate the difficulty of learning. (3) The traversal ensures that each face’s orientation remains consistent within each sub-mesh. Consequently, the generated mesh can be accurately rendered using back face culling, a feature not consistently achieved in prior methods. We will further illustrate these benefits in Section 4.

3.2 AUTO-REGRESSIVE AUTO-ENCODER

Although our decoder is auto-regressive and generates variable-length token sequences, we observe limitations in both generation diversity and its ability to follow conditioning. In contrast, diffusion models, which have been extensively studied (Rombach et al., 2022; Zhang et al., 2024c), demonstrate promising results in addressing these challenges. To apply diffusion models, a key challenge is the requirement of fixed-length data, as mesh generation has a variable-length data structure and the number of triangle faces can vary significantly due to different complexities. Therefore, we propose an **Auto-regressive Auto-encoder (ArAE) to encode the variable-length mesh into a fixed-length latent space**, similar to the role of variational auto-encoders (VAEs) in latent diffusion models. To train our ArAE model, we choose point clouds as the encoder input for geometric information and our tokenized mesh sequences as output. Thus the ArAE model itself also works as a point cloud conditioned mesh generator.

Architecture. The architecture of our ArAE model is illustrated in Figure 2. ArAE consists of a lightweight encoder and an auto-regressive decoder. To extract geometric information from the mesh surface, we follow prior works to sample a point cloud and apply a transformer encoder (Chen et al., 2024d;e; Zhang et al., 2023; 2024c). Specifically, we sample N random points, $\mathbf{X} \in \mathbb{R}^{N \times 3}$, from the surface of the input mesh and use a cross-attention layer to extract the latent code:

$$\mathbf{Z} = \text{CrossAtt}(\mathbf{Q}, \text{PosEmbed}(\mathbf{X})) \quad (1)$$

where $\mathbf{Q} \in \mathbb{R}^{M \times C}$ represents the trainable query embedding with a hidden dimension of C , $\text{PosEmbed}(\cdot)$ is a frequency embedding function for 3D points (Zhang et al., 2023), and $\mathbf{Z} \in \mathbb{R}^{M \times L}$ represents the latent code, $M < N$ and $L < C$ denote the latent size and dimension, respectively. The decoder is an auto-regressive transformer, designed to generate a variable-length mesh token sequence. For simplicity, we adopt the OPT architecture (Zhang et al., 2022), as used in prior works (Siddiqui et al., 2024a; Chen et al., 2024d;b). A learnable embedding converts discrete tokens into continuous features, and a linear head maps predicted features back to classification logits. Stacked causal self-attention layers are employed to predict the next token based on previous tokens. The latent code \mathbf{Z} is prepended to the input before the BOS token, allowing the decoder to learn how to generate a mesh token sequence conditioned on the latent code.

Face Count Condition. Another issue is that given the same input point cloud, multiple plausible meshes with varying numbers of faces and topologies can be generated. The number of faces is particularly crucial as it directly affects the mesh’s complexity (low-poly vs. high-poly) and the generation speed. To manage meshes with a broad range of face counts, we aim to provide some level of explicit control over the targeted number of faces. This control facilitates the estimation of

generation time and the complexity of the generated mesh during inference. We propose a simple face count conditioning method for coarse-grained control. Specifically, we append a learnable face count token after the latent code condition tokens. We bucket face count into different ranges and assign different tokens to each range. For instance, we use four distinct tokens to represent face counts in the following ranges: less than or equal to 1000, between 1000 and 2000, between 2000 and 4000, and greater than 4000. Additionally, during training, we randomly replace these tokens with a fifth unconditional token. This approach ensures that the model still learns to generate meshes without specifying a targeted face count.

Loss Function. The ArAE is trained using the standard cross-entropy loss on the predicted next tokens:

$$\mathcal{L}_{ce} = \text{CrossEntropy}(\hat{\mathbf{S}}[: -1], \mathbf{S}[1 :]) \quad (2)$$

where \mathbf{S} denotes the one-hot ground truth token sequence, and $\hat{\mathbf{S}}$ represents the predicted classification logits sequence. Additionally, to constrain the range of the latent space for easier training of subsequent diffusion models, we apply an L2 norm penalty to the latent code:

$$\mathcal{L}_{reg} = \|\mathbf{Z}\|_2^2 \quad (3)$$

The final loss is a weighted combination of the cross-entropy loss and the regularization term.

3.3 IMAGE-CONDITIONED LATENT DIFFUSION

With the fixed-length latent space provided by our ArAE architecture, it is now feasible to train mesh generation models conditioned on different inputs, akin to how 2D image generation models are trained. Among various input modalities, we showcase the capabilities of our model using single-view images, which are among the most commonly used conditions for mesh generation.

We follow the approach of previous methods (Chen et al., 2023b) and utilize a diffusion transformer (DiT) as the backbone. Specifically, we employ the image encoder from CLIP (Ilharco et al., 2021; Mohammad Khalid et al., 2022) to extract image features for conditioning. Cross-attention layers are used to integrate the image condition into the denoising features, while AdaLN layers incorporate timestep information. We use the DDPM framework (Ho et al., 2020) and mean square error (MSE) loss to train the DiT model. At each training step, we randomly sample a timestep t and a Gaussian noise $\epsilon \in \mathbb{R}^{M \times L}$. The loss is calculated between the predicted noise and the random noise.


4 EXPERIMENTS

4.1 QUALITATIVE RESULTS

Point Cloud Conditioned Generation. We first compare the generated meshes conditioned on point clouds in Figure 5. For the test samples, we use other image-to-3D methods (Zhang et al., 2024c; Li et al., 2024c) to generate dense meshes, ensuring that these samples have never been seen in the training dataset. Notably, MeshAnything (Chen et al., 2024d;e) uses both point coordinates and normal vectors as input to a pretrained encoder (Zhao et al., 2023). In contrast, our encoder takes only point coordinates and is trained from scratch, while still achieving better performance. Our model demonstrates greater stability on challenging test cases, producing more aesthetically pleasing meshes. Additionally, the proposed tokenization algorithm enables the generation of longer sequences, facilitating the modeling of complex shapes and improving the closure of loops for warty surfaces.

Image Conditioned Generation. In Figure 6, we compare the meshes generated from single-view images. Since other auto-regressive methods cannot directly condition on images, we compare our results with a recent optimization-based approach (Wu et al., 2024a). Our end-to-end image-to-mesh generation produces artistic meshes that better capture the semantics of the input images. Additionally, our model exhibits strong generalization on challenging 2D-style or realistic-lighting images, despite being trained exclusively on images rendered from 3D meshes with simple shading.

Face Count Control. In practice, users may have specific requirements for the face count or level of detail in the generated mesh. Our design enables coarse-grained control over the face count by allowing users to specify the face count condition. As demonstrated in Figure 7, our model is capable of generating meshes with different ranges of face count from the same input point cloud.

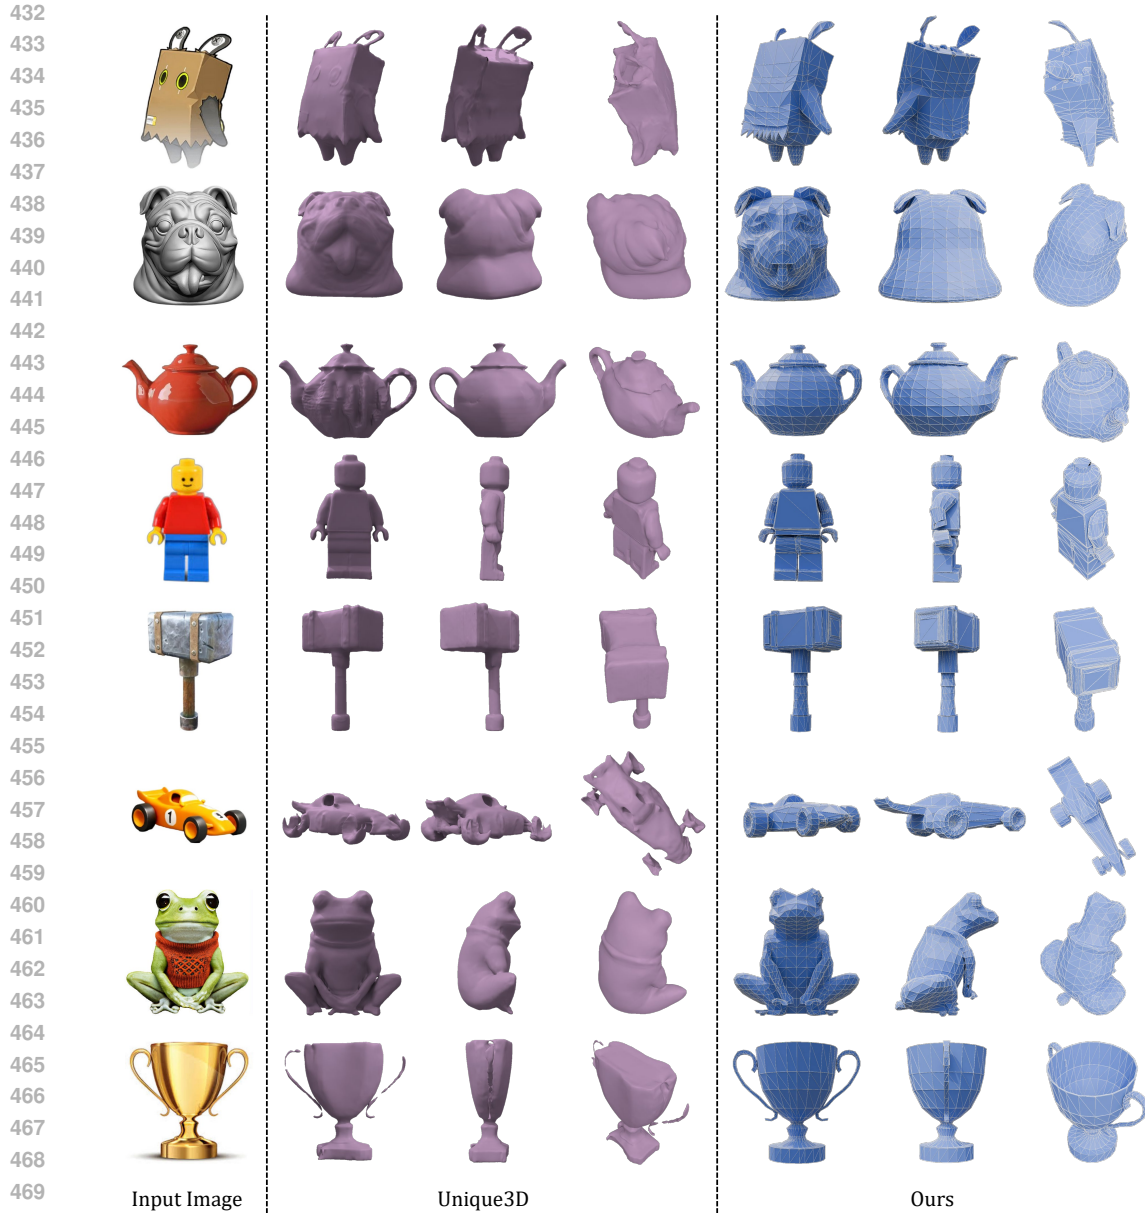


417 **Figure 5: Comparison on point cloud conditioned generation.** We show the reference dense mesh
418 and generated meshes conditioned on randomly sampled point cloud.

421 4.2 QUANTITATIVE RESULTS

422
423 **User Study.** Since evaluation of the mesh aesthetic quality is challenging, we majorly rely on a user
424 study for quantitative comparisons. For a collection of 8 test cases, we render the mesh geometry
425 and wireframe generated from MeshAnything (Chen et al., 2024d), MeshAnythingV2 (Chen et al.,
426 2024e), and our method. Each volunteer is shown 8 samples from mixed random methods, and
427 asked to rate in three aspects: input consistency with the point cloud, triangle aesthetic, and overall
428 mesh quality. We collect results from 20 volunteers with 480 valid scores. As shown in Table 1, our
429 method is preferred across all evaluated aspects.

430 **Inference Speed.** Since we are using next-token prediction, the inference time for each generation is
431 dependent on the length of the sequence. Our model is basically a large transformer, and the major
compute is used for attention calculation. On a A100 GPU with flash-attention (Dao, 2024) and



471 **Figure 6: Comparison on image conditioned generation.** We show the input image and generated
472 meshes. Since Unique3D outputs dense meshes, we only visualize the surface without wireframe.

473
474

	Input Consistency	Triangle Aesthetic	Overall Quality
MeshAnything	2.93	2.43	2.43
MeshAnythingV2	2.64	2.36	2.14
Ours	4.83	4.54	4.58

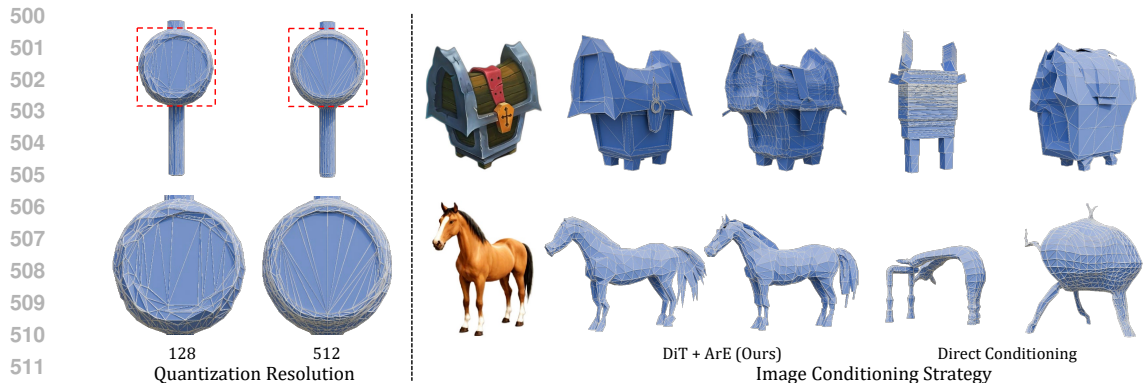
475
476
477
478
479

480 **Table 1: User Study** on point-conditioned mesh generation. The rating is of scale 1-5, the higher
481 the better.

482
483
484 KV cache enabled, our model runs at about 100 tokens per second. Specifically, it takes about 45
485 seconds to generate a mesh with 1,000 faces, about 90 seconds for 2,000 faces, and about 3 minutes
for 4,000 faces.



497 **Figure 7: Coarse-grained face count control.** Our model supports controlling the targeted number
498 of face, which allows to generate meshes with different levels of detail.



513 **Figure 8: Ablation Study** on different quantization resolutions and image conditioning strategies.

516 4.3 ABLATION STUDIES

518 **Quantization Resolution.** We initially experimented with a quantization resolution of 128^3 as
519 used in previous works. While this lower resolution made it easier to train the classification head,
520 we observed noticeable quality degradation, as shown in the left part of Figure 8. To address this
521 problem, we increased the quantization resolution to 512^3 for our final model, which resulted in
522 more accurate vertex positions and smoother surfaces.

523 **Direct Image Conditioning.** We also attempted to directly train an image-conditioned auto-
524 regressive mesh decoder. The image features from the CLIP vision encoder were prepended as
525 condition tokens, and the model directly predicted the mesh token sequence. However, we found
526 that this approach was more difficult to converge and generally exhibited poorer generalization on
527 unseen test cases, as shown in the right part of Figure 8. Consequently, we designed the ArAE model
528 to learn a fixed-length latent space using a simpler point-cloud condition, and we trained latent dif-
529 fusion models to map complex condition modalities to this latent space. This approach led to better
530 generalization and more efficient training.

532 5 CONCLUSION

535 In this work, we propose a novel Auto-regressive Auto-encoder that compresses variable-length tri-
536 angular meshes into fixed-length latent codes, and we further demonstrate the potential of this latent
537 space for image-to-mesh generation. Additionally, we introduce an efficient mesh tokenization al-
538 gorithm that enables scaling up the model, resulting in longer context lengths and higher resolution.
539 Our approach highlights the potential for improving the scalability of auto-regressive models for
practical 3D mesh generation, offering a promising direction for future research and applications.

REFERENCES

- 540
541
542 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,
543 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are
544 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 4
- 545 Ziang Cao, Fangzhou Hong, Tong Wu, Liang Pan, and Ziwei Liu. Large-vocabulary 3d diffusion
546 model with transformer. *arXiv preprint arXiv:2309.07920*, 2023. 3
- 547
548 Anpei Chen, Haofei Xu, Stefano Esposito, Siyu Tang, and Andreas Geiger. Lara: Efficient large-
549 baseline radiance fields. *arXiv preprint arXiv:2407.04699*, 2024a. 3
- 550
551 Hansheng Chen, Jiatao Gu, Anpei Chen, Wei Tian, Zhuowen Tu, Lingjie Liu, and Hao Su. Single-
552 stage diffusion nerf: A unified approach to 3d generation and reconstruction. *arXiv preprint*
553 *arXiv:2304.06714*, 2023a. 3
- 554
555 Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Yue Wu, Zhongdao Wang, James
556 Kwok, Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart- α : Fast training of diffusion transformer
557 for photorealistic text-to-image synthesis, 2023b. 7
- 558
559 Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and
560 appearance for high-quality text-to-3d content creation. *arXiv preprint arXiv:2303.13873*, 2023c.
3
- 561
562 Sijin Chen, Xin Chen, Anqi Pang, Xianfang Zeng, Wei Cheng, Yijun Fu, Fukun Yin, Yanru Wang,
563 Zhibin Wang, Chi Zhang, et al. Meshxl: Neural coordinate field for generative 3d foundation
564 models. *arXiv preprint arXiv:2405.20853*, 2024b. 2, 4, 6
- 565
566 Yiwen Chen, Zilong Chen, Chi Zhang, Feng Wang, Xiaofeng Yang, Yikai Wang, Zhongang Cai, Lei
567 Yang, Huaping Liu, and Guosheng Lin. Gaussianeditor: Swift and controllable 3d editing with
568 gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
Recognition, pp. 21476–21485, 2024c. 3
- 569
570 Yiwen Chen, Tong He, Di Huang, Weicai Ye, Sijin Chen, Jiaxiang Tang, Xin Chen, Zhongang Cai,
571 Lei Yang, Gang Yu, et al. Meshanything: Artist-created mesh generation with autoregressive
572 transformers. *arXiv preprint arXiv:2406.10163*, 2024d. 2, 4, 6, 7, 8, 20
- 573
574 Yiwen Chen, Yikai Wang, Yihao Luo, Zhengyi Wang, Zilong Chen, Jun Zhu, Chi Zhang, and Gu-
575 osheng Lin. Meshanything v2: Artist-created mesh generation with adjacent mesh tokenization,
2024e. URL <https://arxiv.org/abs/2408.02555>. 2, 4, 6, 7, 8, 20, 21
- 576
577 Zhaoxi Chen, Fangzhou Hong, Haiyi Mei, Guangcong Wang, Lei Yang, and Ziwei Liu. Primdiffu-
578 sion: Volumetric primitives diffusion for 3d human generation. *arXiv preprint arXiv:2312.04559*,
2023d. 3
- 579
580 Zhaoxi Chen, Jiaxiang Tang, Yuhao Dong, Ziang Cao, Fangzhou Hong, Yushi Lan, Tengfei Wang,
581 Haozhe Xie, Tong Wu, Shunsuke Saito, et al. 3dtopia-xl: Scaling high-quality 3d asset generation
582 via primitive diffusion. *arXiv preprint arXiv:2409.12957*, 2024f. 3
- 583
584 Zilong Chen, Feng Wang, and Huaping Liu. Text-to-3d using gaussian splatting. *arXiv preprint*
585 *arXiv:2309.16585*, pp. 21401–21412, 2023e. 3
- 586
587 Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander G Schwing, and Liang-Yan Gui. Sdfu-
588 sion: Multimodal 3d shape completion, reconstruction, and generation. In *CVPR*, pp. 4456–4465,
2023. 3
- 589
590 Tri Dao. FlashAttention-2: Faster attention with better parallelism and work partitioning. In *Inter-*
591 *national Conference on Learning Representations (ICLR)*, 2024. 8
- 592
593 Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan
Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. Objaverse-xl: A universe of
10m+ 3d objects. *arXiv preprint arXiv:2307.05663*, 2023a. 3, 18

- 594 Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig
595 Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of anno-
596 tated 3d objects. In *CVPR*, pp. 13142–13153, 2023b. 3, 18
- 597 Michael Garland and Paul S Heckbert. Surface simplification using quadric error metrics. In *Pro-
598 ceedings of the 24th annual conference on Computer graphics and interactive techniques*, pp.
599 209–216, 1997. 19
- 600
- 601 Anchit Gupta, Wenhan Xiong, Yixin Nie, Ian Jones, and Barlas Oğuz. 3dgen: Triplane latent
602 diffusion for textured mesh generation. *arXiv preprint arXiv:2303.05371*, 2023. 3
- 603
- 604 Tiankai Hang, Shuyang Gu, Chen Li, Jianmin Bao, Dong Chen, Han Hu, Xin Geng, and Baining
605 Guo. Efficient diffusion training via min-snr weighting strategy. In *Proceedings of the IEEE/CVF
606 International Conference on Computer Vision*, pp. 7441–7451, 2023. 20
- 607 Zexin He and Tengfei Wang. Openlrm: Open-source large reconstruction models. [https://
608 github.com/3DTopia/OpenLRM](https://github.com/3DTopia/OpenLRM), 2023. 3
- 609
- 610 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 33:
611 6840–6851, 2020. 7
- 612 Fangzhou Hong, Jiuxiang Tang, Ziang Cao, Min Shi, Tong Wu, Zhaoxi Chen, Tengfei Wang, Liang
613 Pan, Dahua Lin, and Ziwei Liu. 3dtopia: Large text-to-3d generation model with hybrid diffusion
614 priors. *arXiv preprint arXiv:2403.02234*, 2024. 3
- 615
- 616 Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli,
617 Trung Bui, and Hao Tan. Lrm: Large reconstruction model for single image to 3d. *arXiv preprint
618 arXiv:2311.04400*, 2023. 2, 3
- 619 Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori,
620 Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hannaneh Hajishirzi, Ali
621 Farhadi, and Ludwig Schmidt. Openclip, July 2021. URL [https://doi.org/10.5281/
622 zenodo.5143773](https://doi.org/10.5281/zenodo.5143773). If you use this software, please cite it as below. 7, 19
- 623
- 624 Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided
625 object generation with dream fields. In *CVPR*, pp. 867–876, 2022. 3
- 626 Heewoo Jun and Alex Nichol. Shap-e: Generating conditional 3d implicit functions. *arXiv preprint
627 arXiv:2305.02463*, 2023. 3
- 628
- 629 Yushi Lan, Fangzhou Hong, Shuai Yang, Shangchen Zhou, Xuyi Meng, Bo Dai, Xingang Pan, and
630 Chen Change Loy. Ln3diff: Scalable latent neural fields diffusion for speedy 3d generation. *arXiv
631 preprint arXiv:2403.12019*, 2024. 3
- 632 Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan
633 Sunkavalli, Greg Shakhnarovich, and Sai Bi. Instant3d: Fast text-to-3d with sparse-view gen-
634 eration and large reconstruction model. *arXiv preprint arXiv:2311.06214*, 2023. 2, 3
- 635
- 636 Mengfei Li, Xiaoxiao Long, Yixun Liang, Weiyu Li, Yuan Liu, Peng Li, Xiaowei Chi, Xingqun
637 Qi, Wei Xue, Wenhan Luo, et al. M-lrm: Multi-view large reconstruction model. *arXiv preprint
638 arXiv:2406.07648*, 2024a. 3
- 639
- 640 Peng Li, Yuan Liu, Xiaoxiao Long, Feihu Zhang, Cheng Lin, Mengfei Li, Xingqun Qi, Shanghang
641 Zhang, Wenhan Luo, Ping Tan, et al. Era3d: High-resolution multiview diffusion using efficient
642 row-wise attention. *arXiv preprint arXiv:2405.11616*, 2024b. 3
- 643
- 644 Weiyu Li, Jiarui Liu, Rui Chen, Yixun Liang, Xuelin Chen, Ping Tan, and Xiaoxiao Long. Crafts-
645 man: High-fidelity mesh generation with 3d native generation and interactive geometry refiner.
arXiv preprint arXiv:2405.14979, 2024c. 2, 3, 7
- 646
- 647 Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten
Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content
creation. In *CVPR*, pp. 300–309, 2023. 2, 3

- 648 Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Zexiang Xu, Hao Su, et al. One-2-3-45:
649 Any single image to 3d mesh in 45 seconds without per-shape optimization. *arXiv preprint*
650 *arXiv:2306.16928*, 2023a. 3
- 651
- 652 Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick.
653 Zero-1-to-3: Zero-shot one image to 3d object. *arXiv preprint arXiv:2303.11328*, 2023b. 2, 3
- 654
- 655 Zhen Liu, Yao Feng, Michael J Black, Derek Nowrouzezahrai, Liam Paull, and Weiyang Liu.
656 Meshdiffusion: Score-based generative 3d mesh modeling. *arXiv preprint arXiv:2303.08133*,
657 2023c. 3
- 658 Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma,
659 Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3d: Single image to 3d
660 using cross-domain diffusion. *arXiv preprint arXiv:2310.15008*, 2023. 3
- 661
- 662 William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction
663 algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pp. 347–353, 1998. 2
- 664 Jonathan Lorraine, Kevin Xie, Xiaohui Zeng, Chen-Hsuan Lin, Towaki Takikawa, Nicholas Sharp,
665 Tsung-Yi Lin, Ming-Yu Liu, Sanja Fidler, and James Lucas. Att3d: Amortized text-to-3d object
666 synthesis. *arXiv preprint arXiv:2306.07349*, 2023. 3
- 667
- 668 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint*
669 *arXiv:1711.05101*, 2017. 18
- 670
- 671 Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Realfusion: 360deg
672 reconstruction of any object from a single image. In *CVPR*, pp. 8446–8455, 2023. 3
- 673
- 674 Oscar Michel, Roi Bar-On, Richard Liu, Sagie Benaim, and Rana Hanocka. Text2mesh: Text-driven
675 neural stylization for meshes. In *CVPR*, pp. 13492–13502, 2022. 3
- 676
- 677 Nasir Mohammad Khalid, Tianhao Xie, Eugene Belilovsky, and Tiberiu Popa. Clip-mesh: Generat-
678 ing textured meshes from text using pretrained image-text models. In *SIGGRAPH Asia*, pp. 1–8,
679 2022. 3, 7
- 680
- 681 Norman Müller, Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Bulo, Peter Kotschieder, and
682 Matthias Nießner. Diffrr: Rendering-guided 3d radiance field diffusion. In *CVPR*, pp. 4328–
683 4338, 2023. 3
- 684
- 685 Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system
686 for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022. 3
- 687
- 688 Evangelos Ntavelis, Aliaksandr Siarohin, Kyle Olszewski, Chaoyang Wang, Luc Van Gool, and
689 Sergey Tulyakov. Autodecoding latent 3d diffusion models. *arXiv preprint arXiv:2307.05445*,
690 2023. 3
- 691
- 692 Yatian Pang, Tanghui Jia, Yujun Shi, Zhenyu Tang, Junwu Zhang, Xinhua Cheng, Xing Zhou, Fran-
693 cis EH Tay, and Li Yuan. Envision3d: One image to 3d with anchor views interpolation. *arXiv*
694 *preprint arXiv:2403.08902*, 2024. 3
- 695
- 696 Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d
697 diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 2, 3
- 698
- 699 Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-
700 Ying Lee, Ivan Skorokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123: One image
701 to high-quality 3d object generation using both 2d and 3d diffusion priors. *arXiv preprint*
arXiv:2306.17843, 2023. 3
- 702
- 703 Lingteng Qiu, Guanying Chen, Xiaodong Gu, Qi Zuo, Mutian Xu, Yushuang Wu, Weihao Yuan,
704 Zilong Dong, Liefeng Bo, and Xiaoguang Han. Richdreamer: A generalizable normal-depth
705 diffusion model for detail richness in text-to-3d. In *Proceedings of the IEEE/CVF Conference on*
Computer Vision and Pattern Recognition, pp. 9914–9925, 2024. 3, 19

- 702 Amit Raj, Srinivas Kaza, Ben Poole, Michael Niemeyer, Nataniel Ruiz, Ben Mildenhall, Shiran
703 Zada, Kfir Aberman, Michael Rubinstein, Jonathan Barron, et al. Dreambooth3d: Subject-driven
704 text-to-3d generation. *arXiv preprint arXiv:2303.13508*, 2023. 3
- 705
706 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
707 resolution image synthesis with latent diffusion models. In *CVPR*, pp. 10684–10695, 2022. 6
- 708 Jarek Rossignac. Edgebreaker: Connectivity compression for triangle meshes. *IEEE transactions*
709 *on visualization and computer graphics*, 5(1):47–61, 1999. 2, 4, 5, 18
- 710
711 Tianchang Shen, Zhaoshuo Li, Marc Law, Matan Atzmon, Sanja Fidler, James Lucas, Jun Gao, and
712 Nicholas Sharp. Spacemesh: A continuous representation for learning manifold surface meshes.
713 In *SIGGRAPH Asia 2024 Conference Papers (SA Conference Papers '24)*, pp. 11, New York, NY,
714 USA, 2024. ACM. doi: 10.1145/3680528.3687634. URL [https://doi.org/10.1145/
715 3680528.3687634](https://doi.org/10.1145/3680528.3687634). 5
- 716 Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. Mvdream: Multi-view
717 diffusion for 3d generation. *arXiv preprint arXiv:2308.16512*, 2023. 3
- 718
719 Yawar Siddiqui, Antonio Alliegro, Alexey Artemov, Tatiana Tommasi, Daniele Sirigatti, Vladislav
720 Rosov, Angela Dai, and Matthias Nießner. Meshgpt: Generating triangle meshes with decoder-
721 only transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
722 *Recognition*, pp. 19615–19625, 2024a. 2, 4, 5, 6
- 723 Yawar Siddiqui, Tom Monnier, Filippos Kokkinos, Mahendra Kariya, Yanir Kleiman, Emilien Gar-
724 reau, Oran Gafni, Natalia Neverova, Andrea Vedaldi, Roman Shapovalov, et al. Meta 3d assetgen:
725 Text-to-mesh generation with high-quality geometry, texture, and pbr materials. *arXiv preprint*
726 *arXiv:2407.02445*, 2024b. 3
- 727 Uriel Singer, Shelly Sheynin, Adam Polyak, Oron Ashual, Iurii Makarov, Filippos Kokkinos, Naman
728 Goyal, Andrea Vedaldi, Devi Parikh, Justin Johnson, et al. Text-to-4d dynamic scene generation.
729 *arXiv preprint arXiv:2301.11280*, 2023. 3
- 730
731 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv*
732 *preprint arXiv:2010.02502*, 2020. 20
- 733
734 Jingxiang Sun, Bo Zhang, Ruizhi Shao, Lizhen Wang, Wen Liu, Zhenda Xie, and Yebin Liu.
735 Dreamcraft3d: Hierarchical 3d generation with bootstrapped diffusion prior. *arXiv preprint*
736 *arXiv:2310.16818*, 2023. 3
- 737
738 Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative
739 gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653*, 2023a. 2, 3
- 740
741 Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. Lgm:
742 Large multi-view gaussian model for high-resolution 3d content creation. *arXiv preprint*
743 *arXiv:2402.05054*, 2024a. 3, 18
- 744
745 Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen.
746 Make-it-3d: High-fidelity 3d creation from a single image with diffusion prior. *arXiv preprint*
747 *arXiv:2303.14184*, 2023b. 3
- 748
749 Shitao Tang, Jiacheng Chen, Dilin Wang, Chengzhou Tang, Fuyang Zhang, Yuchen Fan, Vikas
750 Chandra, Yasutaka Furukawa, and Rakesh Ranjan. Mvdifffusion++: A dense high-resolution
751 multi-view diffusion model for single or sparse-view 3d object reconstruction. *arXiv preprint*
752 *arXiv:2402.12712*, 2024b. 3
- 753
754 Zhicong Tang, Shuyang Gu, Chunyu Wang, Ting Zhang, Jianmin Bao, Dong Chen, and Baining
755 Guo. Volumediffusion: Flexible text-to-3d generation with efficient volumetric encoder. *arXiv*
preprint arXiv:2312.11459, 2023c. 3
- 756
757 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
758 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and
759 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 4

- 756 Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian
757 chaining: Lifting pretrained 2d diffusion models for 3d generation. In *CVPR*, pp. 12619–12629,
758 2023a. 3
- 759 Peng Wang, Hao Tan, Sai Bi, Yinghao Xu, Fujun Luan, Kalyan Sunkavalli, Wenping Wang, Zexi-
760 ang Xu, and Kai Zhang. Pf-lrm: Pose-free large reconstruction model for joint pose and shape
761 prediction. *arXiv preprint arXiv:2311.12024*, 2023b. 2, 3
- 762 Tengfei Wang, Bo Zhang, Ting Zhang, Shuyang Gu, Jianmin Bao, Tadas Baltrusaitis, Jingjing Shen,
763 Dong Chen, Fang Wen, Qifeng Chen, et al. Rodin: A generative model for sculpting 3d digital
764 avatars using diffusion. In *CVPR*, pp. 4563–4573, 2023c. 3
- 765 Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolific-
766 dreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *arXiv*
767 *preprint arXiv:2305.16213*, 2023d. 3
- 770 Zhengyi Wang, Yikai Wang, Yifei Chen, Chendong Xiang, Shuo Chen, Dajiang Yu, Chongxuan Li,
771 Hang Su, and Jun Zhu. Crm: Single image to 3d textured mesh with convolutional reconstruction
772 model. *arXiv preprint arXiv:2403.05034*, 2024. 3
- 773 Xinyue Wei, Kai Zhang, Sai Bi, Hao Tan, Fujun Luan, Valentin Deschaintre, Kalyan Sunkavalli,
774 Hao Su, and Zexiang Xu. Meshlrm: Large reconstruction model for high-quality mesh. *arXiv*
775 *preprint arXiv:2404.12385*, 2024. 3
- 776 Kevin J Weiler. *Topological structures for geometric modeling (Boundary representation, manifold,*
777 *radial edge structure)*. Rensselaer Polytechnic Institute, 1986. 5
- 778 Haohan Weng, Yikai Wang, Tong Zhang, CL Chen, and Jun Zhu. Pivotmesh: Generic 3d mesh
779 generation via pivot vertices guidance. *arXiv preprint arXiv:2405.16890*, 2024a. 2, 4
- 780 Zhenzhen Weng, Jingyuan Liu, Hao Tan, Zhan Xu, Yang Zhou, Serena Yeung-Levy, and Jimei
781 Yang. Single-view 3d human digitalization with large reconstruction models. *arXiv preprint*
782 *arXiv:2401.12175*, 2024b. 2
- 783 Kailu Wu, Fangfu Liu, Zhihan Cai, Runjie Yan, Hanyang Wang, Yating Hu, Yueqi Duan, and
784 Kaisheng Ma. Unique3d: High-quality and efficient 3d mesh generation from a single image,
785 2024a. 3, 7
- 786 Shuang Wu, Youtian Lin, Feihu Zhang, Yifei Zeng, Jingxi Xu, Philip Torr, Xun Cao, and Yao Yao.
787 Direct3d: Scalable image-to-3d generation via 3d latent diffusion transformer. *arXiv preprint*
788 *arXiv:2405.14832*, 2024b. 2, 3
- 789 Desai Xie, Sai Bi, Zhixin Shu, Kai Zhang, Zexiang Xu, Yi Zhou, Sören Pirk, Arie Kaufman, Xin
790 Sun, and Hao Tan. Lrm-zero: Training large reconstruction models with synthesized data. *arXiv*
791 *preprint arXiv:2406.09371*, 2024. 3
- 792 Dejia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Yi Wang, and Zhangyang Wang. Neurallift-360:
793 Lifting an in-the-wild 2d photo to a 3d object with 360deg views. In *CVPR*, pp. 4479–4489,
794 2023a. 3
- 795 Dejia Xu, Ye Yuan, Morteza Mardani, Sifei Liu, Jiaming Song, Zhangyang Wang, and Arash Vahdat.
796 Agg: Amortized generative 3d gaussians for single image to 3d. *arXiv preprint arXiv:2401.04099*,
797 2024a. 3
- 798 Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. Instantmesh:
799 Efficient 3d mesh generation from a single image with sparse-view large reconstruction models.
800 *arXiv preprint arXiv:2404.07191*, 2024b. 3
- 801 Xiang Xu, Joseph Lambourne, Pradeep Jayaraman, Zhengqing Wang, Karl Willis, and Yasutaka
802 Furukawa. Brep-gen: A b-rep generative diffusion model with structured latent geometry. *ACM*
803 *Transactions on Graphics (TOG)*, 43(4):1–14, 2024c. 3

- 810 Yinghao Xu, Hao Tan, Fujun Luan, Sai Bi, Peng Wang, Jiahao Li, Zifan Shi, Kalyan Sunkavalli,
811 Gordon Wetzstein, Zexiang Xu, et al. Dmv3d: Denoising multi-view diffusion using 3d large
812 reconstruction model. *arXiv preprint arXiv:2311.09217*, 2023b. [2](#)
813
- 814 Yinghao Xu, Zifan Shi, Wang Yifan, Hansheng Chen, Ceyuan Yang, Sida Peng, Yujun Shen, and
815 Gordon Wetzstein. Grm: Large gaussian reconstruction model for efficient 3d reconstruction and
816 generation. *arXiv preprint arXiv:2403.14621*, 2024d. [3](#)
- 817 Xingguang Yan, Han-Hung Lee, Ziyu Wan, and Angel X Chang. An object is worth 64x64 pixels:
818 Generating 3d object via image diffusion. *arXiv preprint arXiv:2408.03178*, 2024. [3](#)
819
- 820 Lior Yariv, Omri Puny, Natalia Neverova, Oran Gafni, and Yaron Lipman. Mosaic-sdf for 3d gener-
821 ative models. *arXiv preprint arXiv:2312.09222*, 2023. [3](#)
- 822 Taoran Yi, Jiemin Fang, Guanjun Wu, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Qi Tian, and
823 Xinggang Wang. Gaussiandreamer: Fast generation from text to 3d gaussian splatting with point
824 cloud priors. *arXiv preprint arXiv:2310.08529*, 2023. [3](#)
825
- 826 Xuanyu Yi, Zike Wu, Qihong Shen, Qingshan Xu, Pan Zhou, Joo-Hwee Lim, Shuicheng Yan,
827 Xinchao Wang, and Hanwang Zhang. Mvgamba: Unify 3d content generation as state space
828 sequence modeling. *arXiv preprint arXiv:2406.06367*, 2024. [3](#)
- 829 Biao Zhang, Jiapeng Tang, Matthias Niessner, and Peter Wonka. 3dshape2vecset: A 3d shape rep-
830 resentation for neural fields and generative diffusion models. *arXiv preprint arXiv:2301.11445*,
831 2023. [3](#), [6](#), [18](#)
- 832 Chubin Zhang, Hongliang Song, Yi Wei, Yu Chen, Jiwen Lu, and Yansong Tang. Geolrm:
833 Geometry-aware large reconstruction model for high-quality 3d gaussian generation. *arXiv*
834 *preprint arXiv:2406.15333*, 2024a. [3](#)
835
- 836 Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu.
837 Gs-lrm: Large reconstruction model for 3d gaussian splatting. *arXiv preprint arXiv:2404.19702*,
838 2024b. [3](#)
- 839 Longwen Zhang, Ziyu Wang, Qixuan Zhang, Qiwei Qiu, Anqi Pang, Haoran Jiang, Wei Yang, Lan
840 Xu, and Jingyi Yu. Clay: A controllable large-scale generative model for creating high-quality 3d
841 assets. *arXiv preprint arXiv:2406.13897*, 2024c. [2](#), [3](#), [6](#), [7](#)
842
- 843 Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christo-
844 pher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer
845 language models. *arXiv preprint arXiv:2205.01068*, 2022. [6](#)
- 846 Zibo Zhao, Wen Liu, Xin Chen, Xianfang Zeng, Rui Wang, Pei Cheng, Bin Fu, Tao Chen, Gang Yu,
847 and Shenghua Gao. Michelangelo: Conditional 3d shape generation based on shape-image-text
848 aligned latent representation. *arXiv preprint arXiv:2306.17115*, 2023. [3](#), [7](#)
849
- 850 Xin-Yang Zheng, Hao Pan, Peng-Shuai Wang, Xin Tong, Yang Liu, and Heung-Yeung Shum. Lo-
851 cally attentional sdf diffusion for controllable 3d shape generation. *ACM Transactions on Graph-*
852 *ics (ToG)*, 42(4):1–13, 2023. [3](#)
- 853 Jingyu Zhuang, Chen Wang, Lingjie Liu, Liang Lin, and Guanbin Li. Dreameditor: Text-driven 3d
854 scene editing with neural fields. *arXiv preprint arXiv:2306.13455*, 2023. [3](#)
855
- 856 Zi-Xin Zou, Zhipeng Yu, Yuan-Chen Guo, Yangguang Li, Ding Liang, Yan-Pei Cao, and Song-Hai
857 Zhang. Triplane meets gaussian splatting: Fast and generalizable single-view 3d reconstruction
858 with transformers. *arXiv preprint arXiv:2312.09147*, 2023. [3](#)
859
860
861
862
863

A APPENDIX

A.1 IMPLEMENTATION DETAILS

A.1.1 MESH TOKENIZER

Algorithm 1: Tokenization**Data:** Discretized Vertices $\mathbf{V} = \{\mathbf{v}_i\}_N$, Triangle Faces $\mathbf{T} = \{\mathbf{t}_i\}_M$.**Result:** Array \mathbf{O} to hold the token sequence.

```

873 /* Write a vertex to output. */
874 def WriteVertex (Vertex v):
875     O.append(v.x, v.y, v.z);
876
877 /* Recursive function to compress one face. */
878 def TokenizeFace (HalfEdge c, bool first):
879     c.t.vis = true;
880     if not first:
881         WriteVertex (c.v);
882     if not c.v.vis:
883         c.v.vis = true;
884         O.append(N);
885         TokenizeFace (c.n.o, false);
886     elif c.n.o.t.vis and c.p.o.t.vis:
887         return; /* End of recursion. */
888     elif c.n.o.t.vis:
889         O.append(P);
890         TokenizeFace (c.p.o, false);
891     elif c.p.o.t.vis:
892         O.append(N);
893         TokenizeFace (c.n.o, false);
894     else: /* both left and right triangles are not visited. */
895         O.append(N);
896         TokenizeFace (c.n.o, false);
897         Traverse (c.p.o); /* Begin a new sub-sequence for c.p.o.t. */
898
899 /* Begin a new sub-sequence */
900 def Traverse (HalfEdge c):
901     if not c.t.vis:
902         O.append(B);
903         WriteVertex (c.v);
904         WriteVertex (c.s);
905         WriteVertex (c.e);
906         c.s.vis = 1; c.e.vis = 1;
907         TokenizeFace (c, true);
908
909 /* Initialization */
910 for Vertex v in V:
911     if v is a boundary vertex:
912         v.vis = true;
913 O.append(BOS);
914
915 /* Loop all unvisited faces and traverse. */
916 for Face t in T:
917     if not t.vis:
918         Traverse (t.halfedges[0]);
919 O.append(EOS);

```

Tokenization Algorithm. To formally define the algorithm, we first describe how to use the half-edge notation to refer the mesh data in Figure 9. For a vertex \mathbf{v} , its three coordinates are denoted as

($v.x, v.y, v.z$). For a face t , we store its three half-edges in an array $t.halfedges$. Additionally, we use $v.vis$ and $t.vis$ to indicate whether this vertex or face has already been visited in the algorithm.

Comparison. We compare the original EdgeBreaker algorithm with our modified version and highlight the key differences. For clarity, we use **N** and **P** tokens in the main paper to represent moving to the right or left triangles (the next twin or previous twin half-edges). To remain consistent with the original algorithm, we now revert to using **L** and **R**.

(1) In the original EdgeBreaker algorithm, newly added vertices are stored in a residual form relative to the previous triangle, following the parallelogram rule (Rossignac, 1999). This method can potentially reduce the number of bits when compressed further in binary format. However, since our objective is to convert the mesh into a discrete token sequence rather than binary bits, relative vertex encoding is unnecessary. Moreover, our discretization reduces the accuracy of relative coordinates, so we simply use the absolute coordinates for each vertex.

(2) As illustrated in Figure 10, the original algorithm employs five face type tokens: C, L, E, R, and S. The special S token is used to indicate a bifurcation in the traversal tree. However, since the binary tree must be flattened into a 1D sequence, both branches dependent on the S-type triangle can introduce long-range dependencies, particularly if the first-visited branch is long, as shown in Figure 11. To avoid such dependencies, we explicitly duplicate the necessary vertices before visiting the second branch, which can be interpreted as restarting a sub-mesh sequence. While this approach slightly worsens the compression rate, it effectively eliminates long-range dependencies, which can be difficult for the transformer model to learn.

(3) During tokenization, the C token must be distinguished from the others, but it behaves similarly to the L token during de-tokenization. Since tokenization is unnecessary during inference, we replace all C tokens with L for training. Additionally, we merge the E token into B, as E always follows B after removing the S token. For simplicity, we don't consider special cases such as holes or handles, since they can still be represented at the cost of several degenerated faces, which can be easily cleaned during post-processing. As a result, our modified algorithm uses only three face types: L, R, and B.

A.1.2 AUTO-REGRESSIVE AUTO-ENCODER

Datasets. For the ArAE model, we use meshes from the Objaverse and Objaverse-XL datasets (Deitke et al., 2023b;a). Given that the original datasets include many low-quality 3D assets, we filter the meshes empirically based on metadata captions or file names (Tang et al., 2024a). Our training set comprises approximately 112K meshes with fewer than 4,000 faces from both Objaverse and Objaverse-XL. For finetuning, we select only the higher-quality meshes from Objaverse, totaling around 44K. Each mesh undergoes a preprocessing pipeline where close vertices are merged, collapsed faces are removed, and the mesh is normalized to fit within a unit cube. The vertices are quantized at a resolution of 512^3 . Our tokenizer's vocabulary thus includes 512 coordinate tokens, 3 face type tokens (**N**, **P**, **B**) and 3 special tokens (**BOS**, **EOS**, **PAD**), resulting in a total vocabulary size of 518.

Architecture. We illustrate the detailed architecture of our network in Figure 12. The encoder of our ArAE model employs a cross-attention layer as outlined by (Zhang et al., 2023). We sample 8,192 points uniformly from the mesh surface and use a $4\times$ down-sampled embedding to query the latent code. Consequently, the latent code has a shape of $\mathbf{Z} \in \mathbb{R}^{2048 \times 64}$. The decoder transformer is comprised of 24 self-attention layers, each with a hidden dimension of 1,536 and 16 attention heads. Overall, our ArAE model contains approximately 0.7 billion trainable parameters.

Training and Inference. We train the ArAE model on 64 A100 (80GB) GPUs for approximately one week. The batch size is 4 per GPU, leading to an effective batch size of 256. We use the AdamW optimizer (Loshchilov & Hutter, 2017) with a cosine-decayed learning rate that ranges from 5×10^{-5} to 5×10^{-6} , a weight decay of 0.1, and betas of (0.9, 0.95). Gradient clipping is




Figure 9: Mesh notation starting from half-edge c .

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971




Figure 10: Comparison of the original EdgeBreaker traversal order with our modified version.

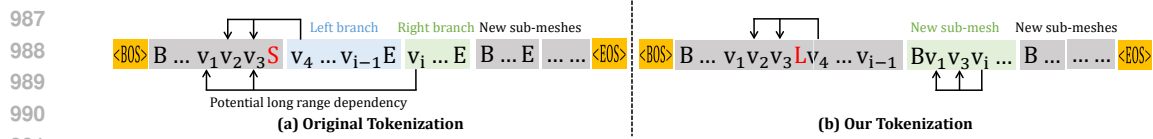


Figure 11: Our tokenization reduces potential long range dependency in the token sequence.

applied with a maximum norm of 1.0. Given that each point cloud input can correspond to multiple plausible meshes but only one ground truth mesh is available, we apply robust data augmentation techniques to address data insufficiency during training: (1) The input mesh is randomly scaled by a factor between $[0.75, 0.95]$. (2) The mesh is randomly rotated along the vertical axis by up to 30 degrees. (3) Random quadric edge collapse decimation (Garland & Heckbert, 1997) is applied to perturb the number of faces. There is a 50% chance at each iteration that the mesh will be decimated to no more than one-quarter of its original number of faces. During inference, we employ a multinomial sampling strategy with a top-k of 10. To ensure that the generated sequence is valid for decoding by the tokenizer, we apply the following rules to post-process the logits of the next token based on the prefix token sequence: (1) The token following **BOS** should be **B**. (2) After **B**, there should be 9 coordinate tokens. (3) After **N** or **P**, there should be 3 coordinate tokens. (4) Otherwise, the next token should be one of **N**, **P**, **B**, or **EOS**.

A.1.3 LATENT DIFFUSION

Datasets. For the image-conditioned diffusion model, we use meshes from the Objaverse dataset along with rendered images from the G-Objaverse dataset (Qiu et al., 2024). The final training dataset comprises approximately 75K 3D meshes, each accompanied by rendered RGB images from 38 different camera poses. We select only 21 camera poses that align with the front face of most objects in Objaverse empirically.

Architecture. Our DiT model consists of 24 layers of self-attention and cross-attention, as detailed in Section 3. Each layer can be expressed as follows:

$$\gamma_1, \beta_1, \alpha_1, \gamma_2, \beta_2, \alpha_2 = \text{Chunk}(\mathbf{T} + \mathbf{t}_e) \quad (4)$$

$$\mathbf{x} = \mathbf{x} + \alpha_1 \times \text{SelfAttention}(\text{LayerNorm}_1(\mathbf{x}) \times (1 + \beta_1) + \gamma_1) \quad (5)$$

$$\mathbf{x} = \mathbf{x} + \text{CrossAttention}(\mathbf{x}, \mathbf{c}) \quad (6)$$

$$\mathbf{x} = \mathbf{x} + \alpha_2 \times \text{FeedForward}(\text{LayerNorm}_2(\mathbf{x}) \times (1 + \beta_2) + \gamma_2) \quad (7)$$

where \mathbf{T} is a layer-wise learnable scale shift table, \mathbf{t}_e denotes the timestep features, and \mathbf{c} denotes the condition features. All attention layers have a hidden dimension of 1,024 and use 16 attention heads. We utilize the pretrained CLIP-H model from OpenCLIP (Ilharco et al., 2021) as the image feature encoder. Features from the last layer of the vision transformer are employed as conditions for the cross-attention layers. The DiT model has approximately 0.5 billion trainable parameters in total.

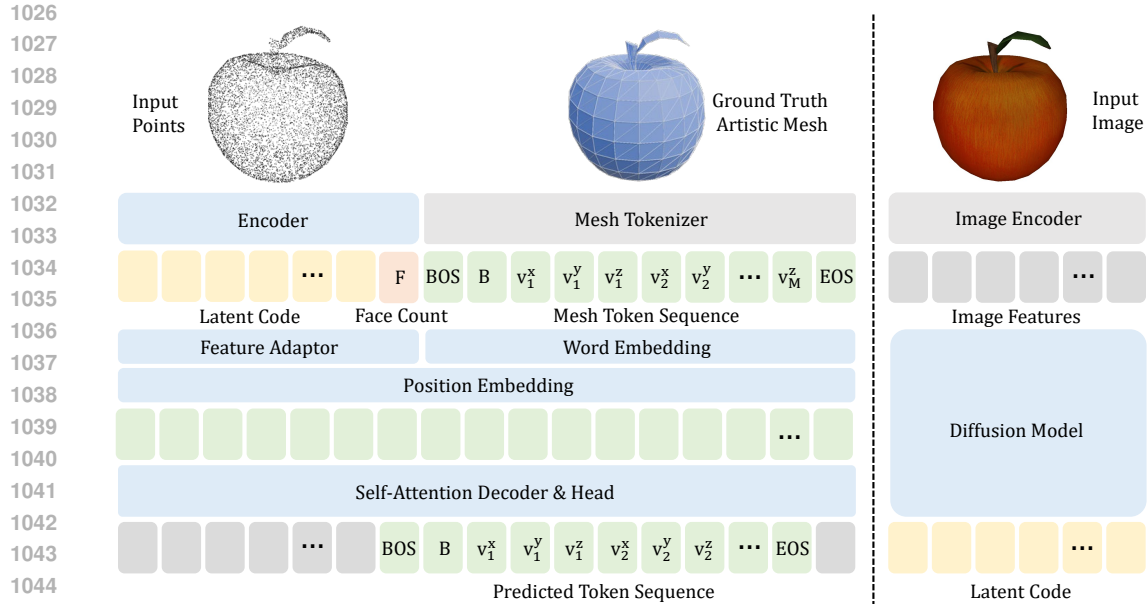


Figure 12: **Network Architecture** of our ArAE and DiT models. Our ArAE contains an encoder that takes a point cloud as input and encodes it to a fixed-length latent code, and an auto-regressive decoder that predicts variable-length mesh token sequence. The latent space can be used to train diffusion models conditioned on other more difficult conditions, such as single-view images.




Figure 13: **Face orientation**. Our model ensures correct face orientation of the generated mesh.

Training and Inference. We train the DiT model on 16 A800 (40GB) GPUs for approximately one week. The lightweight ArAE encoder allows for on-the-fly calculation of the latent code during training, so we can also apply random down-sampling and rotation as data augmentation techniques for the point cloud input. The batch size is set to 32 per GPU, resulting in an effective batch size of 512. We use the same optimizer as for the ArAE model. Training employs the min-SNR strategy (Hang et al., 2023) with a weight of 5.0. The DDPM is configured with 1,000 timesteps and utilizes a scaled linear beta schedule. For classifier-free guidance during inference, 10% of the image conditions are randomly set to zero for unconditional training. During inference, we use the DDIM scheduler (Song et al., 2020) with 100 denoising steps and a Classifier-Free Guidance (CFG) scale of 7.5.

A.2 MORE RESULTS

A.2.1 FACE ORIENTATION

In Figure 13, we show the face orientation and rendering with back-face culling. The tokenization algorithm of our method and MeshAnythingV2 (Chen et al., 2024e) ensures that the face orientation in each mesh sequence is correct, while MeshAnything (Chen et al., 2024d) fails to do this.




Figure 14: Visualization of the generation progress. Please visit our project page for video results.

	Compression Ratio ↓	Sub-sequence Count ↓	Tokenization Speed ↑
AMT	46.2%	199.5	8.6
Ours	47.4%	54.7	25.2

Table 2: **Comparison of the tokenizer algorithm.** Our tokenizer has a comparable compression ratio as AMT (Chen et al., 2024e) but uses much fewer sub-sequences per mesh, which makes model training easier. Moreover, our implementation is three times faster (mesh per second).




Figure 15: **Limitations.** We show some failure cases of our method to illustrate our limitations.

A.2.2 GENERATION PROGRESS

In Figure 14, we show the generation progress by simultaneously de-tokenizing the sequence. Our model successfully learns how to generate a valid sequence for the modified EdgeBreaker algorithm to de-tokenize.

A.2.3 TOKENIZATION ALGORITHM

In Table 2, we compare our mesh tokenization algorithm with AMT from MeshAnythingV2 (Chen et al., 2024e). Both algorithms aim to traverse all triangular faces. While AMT can only move to a fixed side that shares the last two vertices in the sequence, our algorithm allows movement to both the left and right sides. This flexibility reduces the number of sub-sequences, making training easier and improving model robustness, as demonstrated in Figure 5. Although AMT requires only 3 coordinate tokens for a subsequent face and our method requires 4 tokens (with an additional face type token), the amortized compression ratio remains similar.

A.2.4 LIMITATIONS

Despite the promising results, our approach has several limitations: (1) The compression ratio of our tokenization algorithm remains insufficient. Complex meshes, such as those representing game characters, typically require a much higher number of faces, exceeding the current capacity of our

1134 model. (2) Although our model can generate meshes with up to 4,000 faces, the success rate de-
1135 creases as the sequence length increases, and the generation time becomes progressively longer.
1136 (3) While the auto-regressive next-token prediction allows for diverse outputs from the same input,
1137 it also leads to potentially unsatisfactory results. Since there are infinite plausible face topologies
1138 for the same surface, precise control over the generation is difficult to achieve. We illustrate some
1139 failure cases in Figure 15.

1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187