

HiFi-Mesh: High-Fidelity Efficient 3D Mesh Generation via Compact Autoregressive Dependence

Yanfeng Li¹, Tao Tan¹, Qinquan Gao^{2,6}, Zhiwen Cao³, Xiaohong Liu^{4,5*}, Yue Sun^{1*}

¹Intelligent Medical Computing Laboratory, Faculty of Applied Sciences, Macao Polytechnic University

²Fuzhou University

³Sichuan University

⁴Shanghai Jiao Tong University

⁵Shanghai Jiao Tong University Sichuan Research Institute

⁶Imperial Vision Technology Co. Ltd

yuesun@mpu.edu.mo, xiaohongliu@sjtu.edu.cn

Abstract

High-fidelity 3D meshes can be tokenized into one-dimension (1D) sequences and directly modeled using autoregressive approaches for faces and vertices. However, existing methods suffer from insufficient resource utilization, resulting in slow inference and the ability to handle only small-scale sequences, which severely constrains the expressible structural details. We introduce the Latent Autoregressive Network (LANE), which incorporates compact autoregressive dependencies in the generation process, achieving a $6\times$ improvement in maximum generatable sequence length compared to existing methods. To further accelerate inference, we propose the Adaptive Computation Graph Reconfiguration (AdaGraph) strategy, which effectively overcomes the efficiency bottleneck of traditional serial inference through spatiotemporal decoupling in the generation process. Experimental validation demonstrates that LANE achieves superior performance across generation speed, structural detail, and geometric consistency, providing an effective solution for high-quality 3D mesh generation.

1 Introduction

Automatic generation of high-quality 3D meshes is a core task in 3D vision (Sinha et al. 2025). With the rapid development of virtual reality, digital twins, and metaverse applications, there is growing demand for large-scale, accurate, and topologically 3D content generation (Pan et al. 2025; Melnik et al. 2025; Kurai et al. 2025). This makes automatic 3D mesh generation technology practically significant.

In recent years, Large Reconstruction Models (Hong et al. 2023) have enabled large-scale 3D reconstruction by predicting Gaussian splatting. However, the resulting point cloud representations lack topological connectivity, limiting their use in downstream applications. Researchers have thus extended this paradigm to generate 3D meshes with topological structure (Xu et al. 2024; Wu et al. 2024b; Li et al. 2025b). Yet Gaussian splatting-based mesh generation suffers from coarse geometric details and poor topological ac-

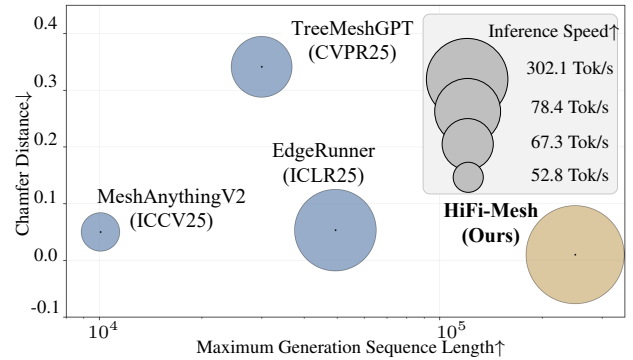


Figure 1: Performance comparisons of EdgeRunner (Tang et al. 2024b), MeshAnythingV2 (Chen et al. 2024b), TreeMeshGPT (Lionar, Liang, and Lee 2025), and HiFi-Mesh (Ours) on Chamfer Distance ↓, Inference Speed ↑, and Maximum Generatable Sequence Length ↑.

curacy due to fundamental differences between volumetric representations and discrete mesh structures, causing precision loss during conversion. Subsequent research has shifted toward directly learning Signed Distance Field (SDF) mesh representations (Zhang et al. 2024; Lai et al. 2025; Wu et al. 2025). However, SDF methods still lose discrete face index information, resulting in poor topological structure with artifacts such as over-smoothing and surface irregularities.

To address this challenge, recent research focuses on tokenizing 3D meshes into 1D sequences, using autoregressive models to generate them for direct modeling of mesh vertices and faces, thereby preserving topological information (Weng et al. 2024). MeshXL (Chen et al. 2024a) directly flattens vertices but results in an inefficient representation, reaching sequence lengths of $3vf$ (f : faces, v : vertices per face). MeshAnythingV2 (Chen et al. 2024c) introduced Adjacent Mesh Tokenization, compressing sequence representation and extending to arbitrary object generation. Although subsequent algorithms have continued to improve upon these foundations, two critical bottlenecks persist.

First, sequence compression algorithms have compres-

*Corresponding author

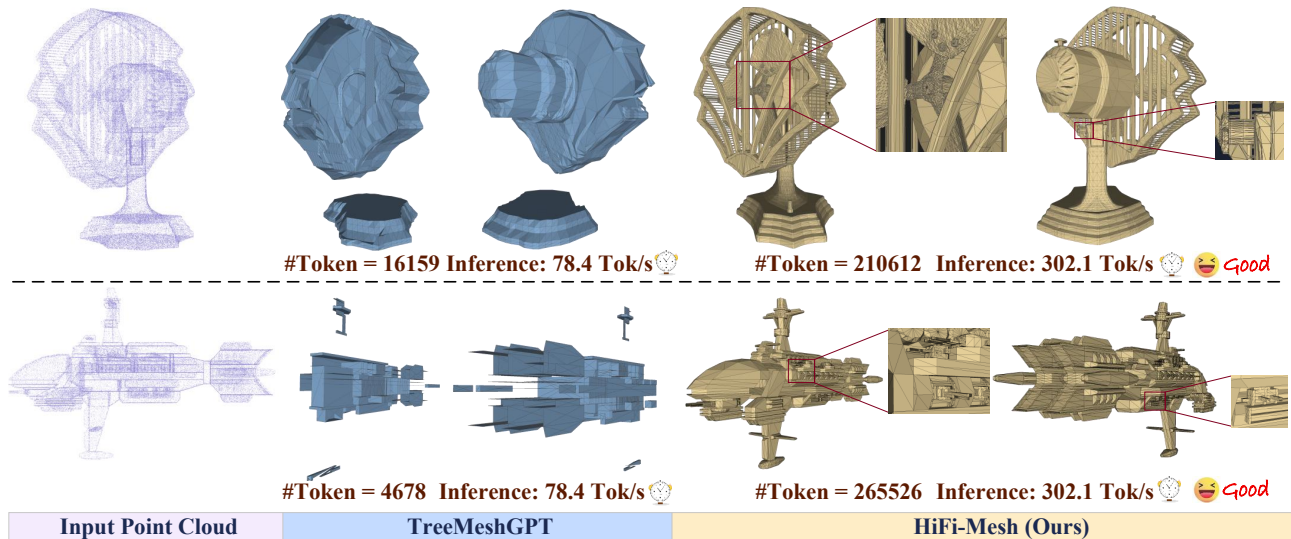


Figure 2: Visual Comparisons between TreeMeshGPT (Lionar, Liang, and Lee 2025) and HiFi-Mesh (Ours). A mesh can be tokenized into a 1D topological sequence. #Token represents sequence length with the density of vertices and faces proportional to it. X Tok/s denotes the inference speed of sequence token. HiFi-Mesh enables detail-rich mesh generation with faster speed.

sion rate limits, and representation of high-fidelity meshes requires sequences on the order of hundreds of thousands. However, constrained by the computational overhead of autoregressive models, current methods are still limited to a maximum generatable length of 50,000 tokens (Tang et al. 2024b). This bottleneck restricts the structural detail of generated meshes. Second, autoregressive model inference must follow strict forward recursive generation from the sequence beginning (Siddiqui et al. 2024). This serialized mechanism significantly constrains inference efficiency. To alleviate this bottleneck, existing methods truncate the context directly, retaining only a limited-length historical sequence as the inference prefix. However, such operations result in substantial information loss, causing significant performance degradation in later stages of long sequence inference.

To address these challenges, we propose HiFi-Mesh for efficient generation of detail-rich 3D meshes, which comprises a Latent Autoregressive Network (LANE) model and Adaptive Computation Graph Reconfiguration (AdaGraph) strategy. The core of the LANE lies in partitioning topological sequences into multiple structural units and introducing sequence semantic reconstruction to build corresponding latent space representations from point cloud as autoregressive dependencies, while ensuring coherence and compactness of structural information. First, the model performs autoregressive generation following subsequence order, with each step depending only on current and preceding latent spaces, reducing long-sequence dependencies. Second, to enhance local modeling capability, we introduce learnable query tokens as structural queriers that dynamically perceive key features in latent spaces, achieving sequential alignment and detail restoration. This design reduces computational overhead while effectively improving maximum generatable length and expressible structural details.

To further enhance inference efficiency, we propose AdaGraph strategy. This strategy leverages the autoregressive dependencies introduced by LANE to achieve spatiotemporal decoupling in generation process. Specifically, AdaGraph first exploits the compactness of latent space representations for rapid hierarchical construction. Subsequently, it decomposes serial inference into multiple independent pathways, where each pathway dynamically reconstructs dedicated computational subgraphs based on target subsequences and selectively activates corresponding latent spaces for parallel generation. Compared to previous methods, AdaGraph maintains complete contextual information, avoiding performance degradation while achieving approximately 300% speed improvement. Figures 1 and 2 shows performance and visual comparisons, respectively.

Our contributions are as follows:

- We propose HiFi-Mesh, which incorporates compact autoregressive dependencies in sequence generation to enable rapid generation of detail-rich 3D meshes.
- LANE and AdaGraph are introduced to enhance maximum generatable sequence length and enable parallel inference, respectively.
- Compared to existing methods, HiFi-Mesh demonstrates superior performance in generation speed, structural detail, and geometric consistency with point clouds, providing valuable insights for addressing such challenges.

2 Related Work

2.1 3D Generation

Traditional rule-based modeling techniques (Wu et al. 2024a) have gradually been replaced by diverse learned representations, including voxel generation (Zhao et al. 2025b; Chen et al. 2025), implicit function encoding (Lim and Lee

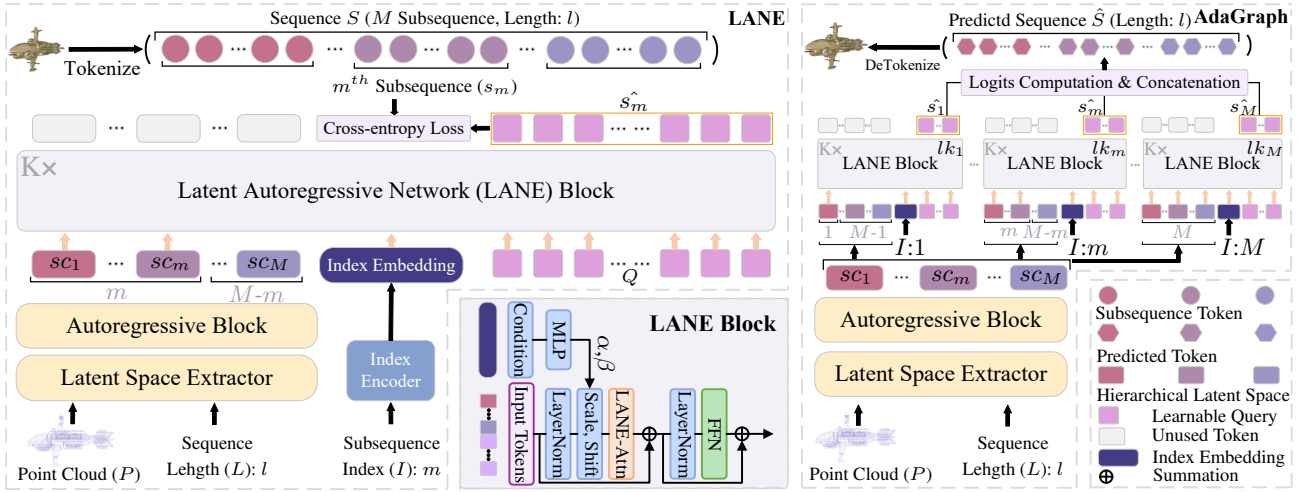


Figure 3: The illustration of our proposed method. During training, the purpose of LANE model (*left*) uniformly divides the topological sequence into M subsequences for sequential generation. The model takes point cloud P and sequence length L as inputs, employing a Latent Space Extractor and Autoregressive Block for sequence semantic reconstruction to build hierarchical latent space representations $\{sc_m\}_{m \in [1, M]}$ as compact autoregressive dependencies. Subsequently, K LANE Blocks serve as generation components, utilizing Learnable Queries Q as latent space queriers. Guided by subsequence index $I = \{m\}_{m \in [1, M]}$, each step generates the m^{th} subsequence \hat{s}_m from the only first m latent spaces. During inference, the AdaGraph strategy (*right*) rapidly generates hierarchical latent space representations $\{sc_m\}_{m \in [1, M]}$ and instantiates M independent pathways $\{lk_m\}_{m \in [1, M]}$ composed of LANE Blocks. Each pathway dynamically reconstructs a dedicated computational subgraph based on the target subsequence and selectively activates corresponding latent spaces, enabling fast parallel generation. Finally, mesh generation is completed through subsequence concatenation and detokenization operations.

2024; Di et al. 2025), and direct mesh construction methods (Zhou et al. 2024; Wu et al. 2024c). Regarding model architectures, Generative Adversarial Networks (Wang, Lan, and Li 2024) are giving way to diffusion models (Li et al. 2025c,a) and specialized end-to-end frameworks (Tang et al. 2024a), which exhibit superior capabilities in generating 3D assets. These advancements provide a more efficient methodological foundation for 3D content creation.

2.2 3D Mesh Generation

Meshes precisely describe object surface geometry (Liu et al. 2024). Early research used Gaussian splatting to predict 3D meshes from 2D images (Hu et al. 2025; Huang et al. 2025), but 2D images lack depth and geometric details, creating an information bottleneck. To address this, researchers generate SDF mesh representations directly (Zhang et al. 2023; Hunyuan3D et al. 2025; Xiang et al. 2025). While these methods reduce information loss, their reliance on fixed thresholds causes geometric discontinuities. Recent studies tokenize 3D meshes into 1D sequences, using autoregressive models to directly model mesh vertices and faces. Methods like VAT and MeshAnything use VQVAE for tokenization (Zhang, Xiong, and Xu 2024; Chen et al. 2024b), but suffer from over-compression and significant geometric detail loss. Alternatively, MeshArt and Mesh-RFT directly tokenize meshes into topological sequences with specific ordering (Gao et al. 2025; Liu et al. 2025), preserving complete geometric and topological information. However,

while achieving high information fidelity, the massive sequence scale for detail-rich meshes creates prohibitive computational overhead for autoregressive models.

2.3 Autoregressive Mesh Generation

When processing detail-rich meshes, significant computational efficiency challenges arise. While BPT and EdgeRunner reduced sequence lengths through improved tokenization (Weng et al. 2025; Tang et al. 2024b), these meshes still require hundreds of thousands of tokens, exceeding current computational limits. Meshtron (Hao et al. 2024) addressed this with sequence fragment-based adjacency structure prediction. Although this significantly increased the scale of generatable meshes, it further reduced inference speed, which was already constrained by forward recursive generation mechanisms, showing approximately 30% slower compared to TreeMeshGPT (Lionar, Liang, and Lee 2025). In this work, we propose a novel framework that reduces computational overhead to handle longer sequences. Crucially, we further overcome inference bottlenecks, enabling rapid generation of detail-rich 3D meshes.

3 Method

3.1 Latent Autoregressive Network

The Latent Autoregressive Network (LANE) model (Figure 3, *left*) uniformly splits the topological sequence S into M consecutive subsequences $\{s_m\}_{m \in [1, M]}$ and performs sequential generation conditioned on subsequence indexes

$I = \{m\}_{m \in [1, M]}$, sequence length L , and point cloud P . The core principle of model is to accurately reconstruct local geometric structures into corresponding latent spaces while ensuring structural coherence. The model then uses compact latent space representations to replace the long historical token sequences required for prediction, thereby reducing computational cost, improving structural detail expressiveness. Specifically, the model consists of three components: Latent Space Extractor, Autoregressive Block, LANE block.

Latent Space Extractor. The Latent Space Extractor employs a Point Cloud Encoder to extract highly compact structural features from high-resolution input point clouds P , then reconstructs fine-grained features through learnable Upsample modules that gradually decode the compressed information. Compared to previous methods, our input point clouds have 16 times higher resolution, containing richer geometric details that establish a foundation for enhanced structural detail expressiveness. Specifically, we perform four random samplings, $\{X_i\}_{i \in [1, 4]}$, from the input mesh surface, containing $\{N_i\}_{i \in [1, 4]}$ points respectively, where $X_i \in R^{N_i \times 3}$, and $N_2 < N_3 < N_4 < N_1$. X_1 serves as the input point cloud P and is encoded into a latent code by the Point Cloud Encoder built from cross-attention, using X_2 as query. Subsequently, X_3 and X_4 act as queries in cross-attention-based Upsample modules to progressively obtain a more refined latent code Z .

Additionally, this Extractor includes a Latent Space Constructor that first initializes corresponding latent spaces in subsequence order, then leverages the sequence length L to guide the latent spaces in accurately partitioning subsequence structure boundaries in the latent code, achieving precise extraction and encoding of target features. The construction process of m^{th} latent space is defined as:

$$sc_m^e = CrossAttn([sc_m^{init}; L_e], Z), \quad (1)$$

where, $CrossAttn(\cdot, \cdot)$ denotes the cross-attention layer with the left side as Query. sc_m^{init} represents the m^{th} initialized latent space, where $m \in [1, M]$. L_e and Z denote the embedding of sequence length L and the latent code reconstructed by the Upsample modules, respectively. $[\ ; \]$ indicates vertical concatenation of multiple information sources. Detailed designs are presented in technical appendix.

Autoregressive Block. The core of this component is to transform the encoded latent spaces $\{sc_m^e\}_{m \in [1, M]}$ into hierarchical latent space representations $\{sc_m\}_{m \in [1, M]}$ to capture dependencies between subsequences and maintain global consistency of topological structure. Specifically, this component utilizes an autoregressive mechanism to sequentially process $\{sc_m^e\}_{m \in [1, M]}$ generated by the Latent Space Extractor, thereby constructing $\{sc_m\}_{m \in [1, M]}$ that reflects the hierarchical nature of the overall topological structure. This representation preserves local detail information while ensuring global coherence, providing a high-quality feature representation foundation for subsequent sequence generation tasks. Its expression is as follows:

$$[sc_1; sc_2; \dots; sc_M] = CausalAttn([sc_1^e; sc_2^e; \dots; sc_M^e]), \quad (2)$$

where, $CausalAttn(\cdot)$ represents the causal attention mechanism in autoregressive models, $[\ ; \dots \]$ denotes vertical concatenation of multiple information sources.

LANE Block. The LANE Block, denoted as \mathcal{E}_{block} , aims to sequentially generate subsequences from the constructed hierarchical latent space representations $\{sc_m\}_{m \in [1, M]}$. Specifically, learnable query tokens Q are introduced as queriers. When generating the m^{th} subsequence, coherent structural features are extracted from the latent spaces $\{sc_1, sc_2, \dots, sc_{m-1}\}$ corresponding to preceding subsequences, thereby eliminating dependence on long historical token sequences. Subsequently, with the aid of subsequence index I , Q further perceives structural details from the current corresponding latent space $\{sc_m\}$, enhancing local modeling capability. The interaction can be defined as:

$$V_{out} = \mathcal{E}_{block}(\{sc_1, sc_2, \dots, sc_m\}, I_m^e, Q) \quad (3)$$

where, I_m^e represents the embedding of subsequence index when I takes the value m . The forward pass of \mathcal{E}_{block} is illustrated in Figure 3, where the specific design and definition of LANE-Attn to utilize $\{sc_m\}_{m \in [1, M]}$ will be presented in technical appendix. This block is repeated K times.

Loss Function. LANE model, denoted as \mathcal{E} , employs cross-entropy loss to predict m^{th} subsequence tokens:

$$\begin{aligned} \mathcal{L}_{ce} &= CrossEntropy(\hat{s}_m, s_m), \\ \hat{s}_m &= \mathcal{E}(P, L, I = m), \end{aligned} \quad (4)$$

where \hat{s}_m and s_m denotes the predicted classification logits subsequence and the ground truth subsequence, respectively.

3.2 Adaptive Computation Graph Reconfiguration

The AdaGraph strategy (Figure 3, *right*), leverages the compact autoregressive dependencies introduced by LANE model to achieve spatiotemporal decoupling in generation process, thereby further improving inference speed. Specifically, this strategy comprises two core stages: Hierarchical Latent Space Construction and Subgraph Reconstruction.

Hierarchical Latent Space Construction. Using the Latent Space Extractor and Autoregressive Block, the hierarchical latent space representations are rapidly constructed from the input point cloud P under the guidance of sequence length L . This process is denoted as \mathcal{T} and defined as:

$$[sc_1; sc_2; \dots; sc_M] = \mathcal{T}(P, L). \quad (5)$$

This process follows the definitions in Equations 1 and 2.

Subgraph Reconstruction. LANE internally instantiates M independent pathways $\{lk_m\}_{m \in [1, M]}$, with each pathway responsible for generating specific subsequences. Each pathway dynamically reconstructs its dedicated computational subgraph and selectively activates corresponding latent spaces, achieving spatiotemporal decoupling in the generation process. This strategy employs batch processing to coordinate multi-pathway collaboration, leveraging the advantages of modern parallel computing architectures to enable efficient parallel generation, significantly improving

Methods/Metrics	Generation Metric				Efficiency Metric		
	Chamfer ↓ ($\times 10^{-1}$)	Normal ↓ ($\times 10^{-1}$)	Point-to-Mesh		MOS ↑ %	Length ↑ #Token	Inference ↑ Tok/s
			Mean ↓ ($\times 10^{-2}$)	Hausorff ↓ ($\times 10^{-2}$)			
TreeMeshGPT (Lionar, Liang, and Lee 2025)	3.413	2.001	3.242	8.253	5.26	30,000	78.4
MeshAnythingV2 (Chen et al. 2024b)	<u>0.501</u>	4.951	2.631	6.374	3.72	10,083	67.3
EdgeRunner (Tang et al. 2024b)	0.532	6.339	<u>0.921</u>	<u>3.764</u>	<u>9.31</u>	<u>49,406</u>	52.8
HiFi-Mesh (Ours)	0.075	<u>3.621</u>	0.638	2.251	81.17	300,000	302.1

Table 1: Quantitative comparisons on seven metrics. **Bold** and underlined indicate the best and second-best performance, respectively. Despite minor Normal Consistency metric limitations, HiFi-Mesh excels in mesh consistency and inference speed. User studies confirm the aesthetic and detailed quality of generated meshes. Around 300% boost in inference speed enables processing $6\times$ longer sequences, achieving high-fidelity 3D mesh generation. All test results are obtained using a single H20.

overall inference speed. This process is represented as:

$$\begin{aligned} \hat{S} &= [\hat{s}_1; \dots; \hat{s}_m; \dots; \hat{s}_M], \\ \hat{s}_m &= lk_m(\{sc_1, sc_2, \dots, sc_m\}, I = m, Q), \end{aligned} \quad (6)$$

where each pathway consists of K LANE Blocks, and the generation process follows Equation 3. Finally, by concatenating all predicted classification logic subsequences $\{\hat{s}_m\}_{m \in [1, M]}$ and performing detokenization operation on \hat{S} , the 3D mesh is generated.

4 Experiment

4.1 Implementation

Training Details. The experiments are implemented using PyTorch, conducted on 8 NVIDIA H20 GPUs. We employ the Adam optimizer and optimize for six weeks, with a learning rate decay from $1e-4$ to $1e-6$. Flash Attention (Shah et al. 2024) is integrated to optimize computation. All the results are tested on a single H20 GPU. Model training utilize the Objaverse dataset (Deitke et al. 2023; Tang et al. 2024a).

Tokenization Algorithm. This paper leverages tokenization algorithm of EdgeRunner (Tang et al. 2024b) for autoregressive model compatibility, optimizing the balance between geometric fidelity and compression ratio. The topological sequence labeling is constructed by quantizing coordinate outputs into 512 discrete categories.

4.2 Comparisons

Three algorithms are selected as comparisons: MeshAnythingV2 (Chen et al. 2024b), EdgeRunner (Tang et al. 2024b), and TreeMeshGPT (Lionar, Liang, and Lee 2025). Four generative metrics are employed to evaluate mesh generation quality: Chamfer Distance (CD), Normal Consistency (NC), Point-to-Mesh Distance (PMD), and MOS. Two efficiency metrics are used to assess model efficiency: Maximum Generateable Sequence Length (ML), and Inference Speed (IS). CD evaluates spatial proximity between reconstructed and original mesh. NC assesses the consistency of normal vectors between adjacent faces. PMD measures the distance from points in the point cloud to the reconstructed mesh surface, evaluated from both Mean distance and Hausdorff distance perspectives. MOS reflects the proportion of

participants favoring a specific model. To collect objective evaluations, ten users are recruited for the assessment task. Length indicates the maximum sequence length that the model can generate. Inference represents the number of token that can be generated per second. The generative metrics are implemented by the PyTorch3D (Ravi et al. 2020).

Quantitative Comparisons. The complete comparison results are shown in Table 1. In terms of generation quality, except for a slight deficiency in Normal Consistency, HiFi-Mesh demonstrates superior performance in mesh consistency and robustness (Chamfer Distance and Point-to-Mesh Distance). User preference studies further validate the aesthetic value and expressible structural detail of our generated meshes. Regarding efficiency of resource utilization, we enable a $6\times$ increase in maximum generatable sequence length and further leverage the AdaGraph strategy to achieve a 300% improvement approximately in inference speed, allowing rapid generation of detail-rich 3D meshes.

Qualitative Comparisons. Figure 4 shows visual comparisons between HiFi-Mesh and three baseline methods (Lionar, Liang, and Lee 2025; Tang et al. 2024b; Chen et al. 2024b). Sequence generation scale directly affects mesh reconstruction capability. MeshAnythingV2 suffers from severe limitations in maximum generatable sequence scale, leading to most significant structural information loss. While EdgeRunner shows theoretical improvements in generation length, it still exhibits a strong preference for small-scale sequences in practical applications. TreeMeshGPT achieves relatively balanced performance and can generate larger sequences, resulting in some improvement in mesh details. However, the maximum sequence generation capabilities of these three methods remain insufficient to fully capture the geometric features of complex surfaces, limiting mesh structural detail expressiveness and generation stability. In contrast, HiFi-Mesh significantly increases the upper limit of sequence generation scale, enabling the model to more effectively learn and preserve surface detail information, thereby achieving stable generation of detail-rich 3D meshes.

Functional Presentation. As shown in Figure 5, introducing surface damage to input meshes during training not only enhances model robustness but also enables HiFi-Mesh to perform mesh repair tasks: sampling point clouds from

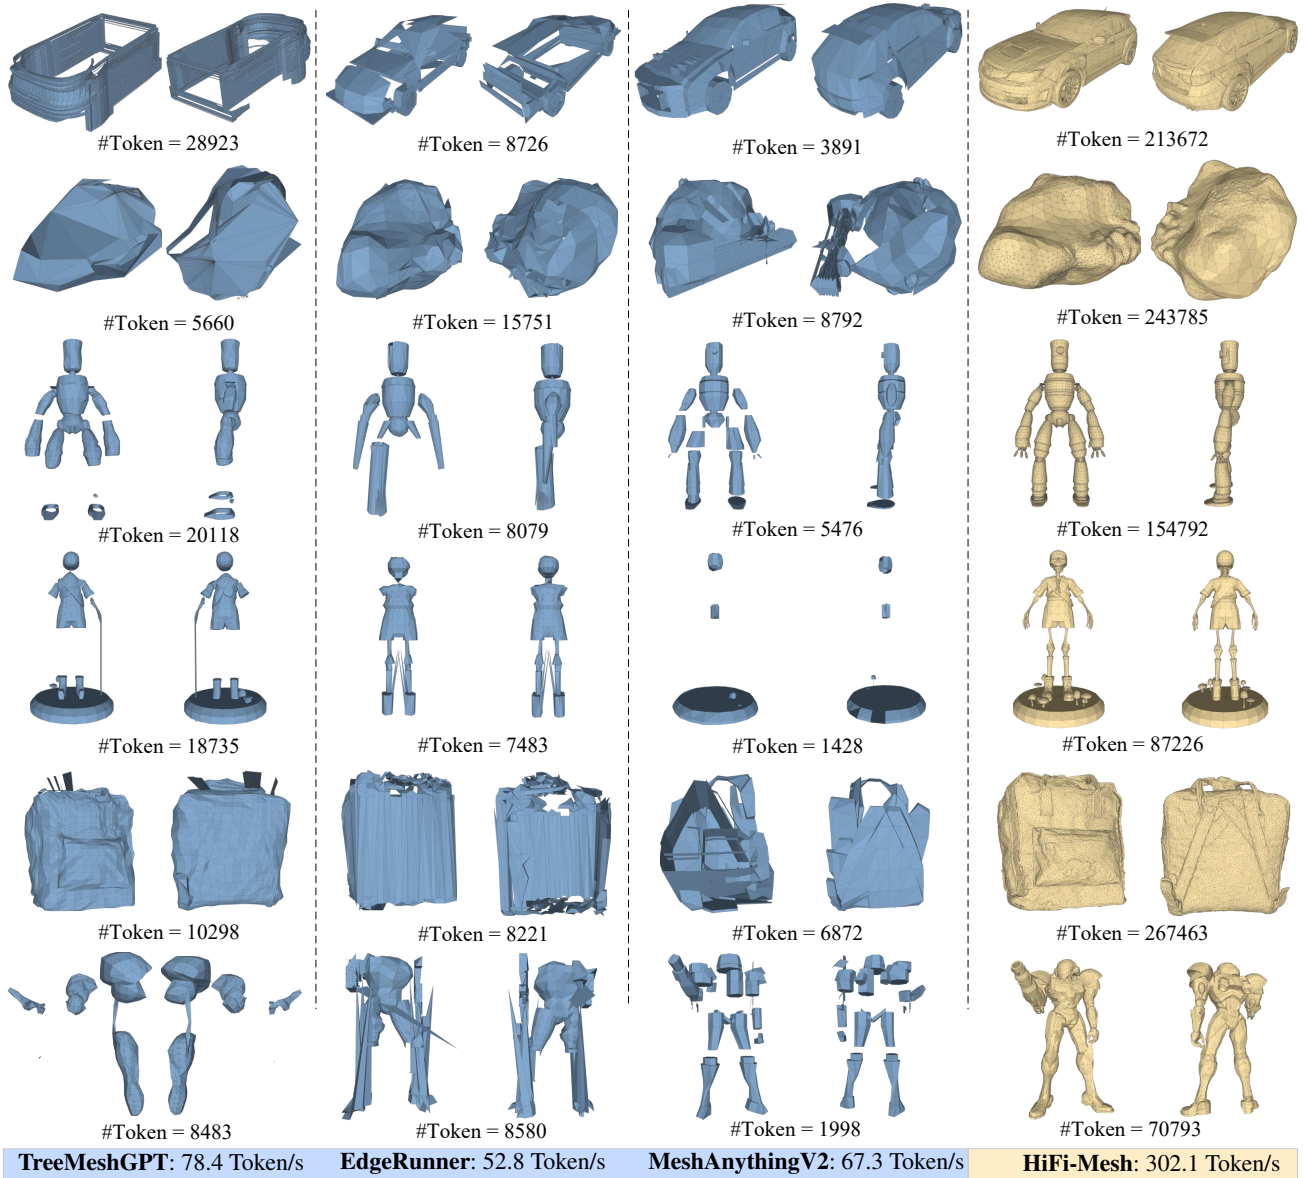


Figure 4: Qualitative comparisons between TreeMeshGPT (Lionar, Liang, and Lee 2025), EdgeRunner (Tang et al. 2024b), MeshAnythingV2 (Chen et al. 2024b) and HiFi-Mesh (Ours). MeshAnythingV2 exhibits limitations in maximum sequence scale, causing most significant structural information loss. EdgeRunner shows theoretical improvements but maintains preference for small-scale sequences. TreeMeshGPT achieves some improvement in detail modeling but still suffers from insufficient sequence generation capability. Finally, HiFi-Mesh enables superior preservation of details and generation of detail-rich meshes.

defective meshes and then reconstructing detailed, high-quality meshes. Additionally, by incorporating the sequence length L as a control parameter, HiFi-Mesh achieves length control over the generation process.

4.3 Ablation Study

Components Effectiveness. As shown in Figure 6. (w/o) $\{s_{C_m}\}_{[1,M]}$ denotes the directly use of global geometric features extracted by the Point Cloud Encoder to generate target subsequences without constructing hierarchical latent space

representations in HiFi-Mesh. Experimental results show that the absence of coherent prior structural features and local detail information provided by latent spaces leads to local topological breaks, geometric integrity loss, and inability to express detailed structures in the generated results.

Table 2 quantifies the importance of hierarchical latent space representations and AdaGraph strategy. Hierarchical latent spaces effectively capture local detail features while maintaining global geometric coherence, significantly improving the generation performance of the model. Addition-




Figure 5: Functional presentation. Point clouds are sampled from imperfect meshes. HiFi-Mesh generates meshes with rich details guided by different sequence lengths, while EdgeRunner (Tang et al. 2024b) faces local topological breaks.




Figure 6: Ablation study. (w/o) $\{sc_m\}_{[1,M]}$ represents without constructing hierarchical latent space representations in HiFi-Mesh. It leads to local topological breaks, geometric integrity loss, and inability to express detailed structures.

$\{sc_m\}_{[1,M]}$	Chamfer \downarrow (10^{-1})	Normal \downarrow (10^{-1})	AdaGraph	Inference \uparrow Tok/s
✗	1.291	6.339	✗	121.7
✓	0.157	1.972	✓	302.1

Table 2: Ablation study. $\{sc_m\}_{[1,M]}$ represents the construction of hierarchical latent space representations in our method HiFi-Mesh. The hierarchical latent spaces provide compact autoregressive dependencies, significantly improving the generative performance of HiFi-Mesh. Additionally, the AdaGraph strategy is introduced to further achieve approximately $3\times$ inference acceleration.

ally, the AdaGraph strategy leverages modern parallel computing architectures to further achieve approximately $3\times$ inference acceleration. Together, they validate the dual advantages of the proposed method in both efficiency and quality. More experiments are presented in technical appendix.

Resource Consumption. Figure 7 presents a comparative analysis of resource consumption between HiFi-Mesh and traditional prediction methods of leveraging long historical sequence during training. The experiment uses a batch size of 4, with EdgeRunner (Tang et al. 2024b) as the traditional baseline, evaluating computational complexity and memory consumption of both approaches across different maximum generatable sequence lengths. Results show that HiFi-Mesh does not demonstrate efficiency advantages over EdgeRunner for short sequences, primarily due to additional computational overhead from constructing hierarchical latent space representations. However, when maximum sequence length exceeds 5,000, HiFi-Mesh begins showing significant efficiency advantages that become increasingly pronounced with longer sequences, validating superior efficiency of HiFi-Mesh in long sequence processing tasks.




Figure 7: Memory consumption and computational complexity for varying maximum generatable sequence lengths. HiFi-Mesh demonstrates gains when length exceeds 5K, highlighting superior efficiency for sequence processing.

5 Conclusion

This paper proposes HiFi-Mesh, a method for rapidly generating detail-rich 3D meshes. The approach leverages compact coherent structural representations and local detail information in latent spaces as autoregressive dependencies, effectively reducing reliance on long historical token sequences and significantly lowering computational overhead. This improvement enables longer maximum generatable sequences, achieving higher structural details. Additionally, we introduce the AdaGraph strategy, which decouples the generation process spatiotemporally into multiple parallel sub-sequences, substantially improving inference speed. Experimental results demonstrate that HiFi-Mesh achieves superior performance in generation speed, structural details, and geometric consistency with input point clouds. This provides an insightful solution to long sequence generation problems and marks an important step toward automatically generating artist-friendly 3D meshes.

Acknowledgements

This work was supported by the grant from Science and Technology Development Fund of Macao (0009/2025/ITP1). This work is supported by Macao Polytechnic University Grant (RP/FCA-08/2024). The work was supported in part by the National Natural Science Foundation of China under Grant 62301310, and in part by Sichuan Science and Technology Program under Grant 2024NS-FSC1426.

References

- Chen, S.; Chen, X.; Pang, A.; Zeng, X.; Cheng, W.; Fu, Y.; Yin, F.; Wang, B.; Yu, J.; Yu, G.; et al. 2024a. Meshxl: Neural coordinate field for generative 3d foundation models. *Advances in Neural Information Processing Systems*, 37: 97141–97166.
- Chen, Y.; He, T.; Huang, D.; Ye, W.; Chen, S.; Tang, J.; Chen, X.; Cai, Z.; Yang, L.; Yu, G.; et al. 2024b. Meshanything: Artist-created mesh generation with autoregressive transformers. *arXiv preprint arXiv:2406.10163*.
- Chen, Y.; Wang, Y.; Luo, Y.; Wang, Z.; Chen, Z.; Zhu, J.; Zhang, C.; and Lin, G. 2024c. Meshanything v2: Artist-created mesh generation with adjacent mesh tokenization. *arXiv preprint arXiv:2408.02555*.
- Chen, Z.; Tang, J.; Dong, Y.; Cao, Z.; Hong, F.; Lan, Y.; Wang, T.; Xie, H.; Wu, T.; Saito, S.; et al. 2025. 3dtopia-xl: Scaling high-quality 3d asset generation via primitive diffusion. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 26576–26586.
- Deitke, M.; Liu, R.; Wallingford, M.; Ngo, H.; Michel, O.; Kusupati, A.; Fan, A.; Laforte, C.; Voleti, V.; Gadre, S. Y.; VanderBilt, E.; Kembhavi, A.; Vondrick, C.; Gkioxari, G.; Ehsani, K.; Schmidt, L.; and Farhadi, A. 2023. Objaverse-XL: A Universe of 10M+ 3D Objects. *arXiv preprint arXiv:2307.05663*.
- Di, D.; Yang, J.; Luo, C.; Xue, Z.; Chen, W.; Yang, X.; and Gao, Y. 2025. Hyper-3dg: Text-to-3d gaussian generation via hypergraph. *International Journal of Computer Vision*, 133(5): 2886–2909.
- Gao, D.; Siddiqui, Y.; Li, L.; and Dai, A. 2025. MeshArt: Generating Articulated Meshes with Structure-guided Transformers. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 618–627.
- Hao, Z.; Romero, D. W.; Lin, T.-Y.; and Liu, M.-Y. 2024. Meshtron: High-fidelity, artist-like 3d mesh generation at scale. *arXiv preprint arXiv:2412.09548*.
- Hong, Y.; Zhang, K.; Gu, J.; Bi, S.; Zhou, Y.; Liu, D.; Liu, F.; Sunkavalli, K.; Bui, T.; and Tan, H. 2023. Lrm: Large reconstruction model for single image to 3d. *arXiv preprint arXiv:2311.04400*.
- Hu, H.; Yin, T.; Luan, F.; Hu, Y.; Tan, H.; Xu, Z.; Bi, S.; Tulsiani, S.; and Zhang, K. 2025. Turbo3d: Ultra-fast text-to-3d generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 23668–23678.
- Huang, Z.; Boss, M.; Vasishta, A.; Rehg, J. M.; and Jampani, V. 2025. Spar3d: Stable point-aware reconstruction of 3d objects from single images. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 16860–16870.
- Hunyuan3D, T.; Yang, S.; Yang, M.; Feng, Y.; Huang, X.; Zhang, S.; He, Z.; Luo, D.; Liu, H.; Zhao, Y.; et al. 2025. Hunyuan3D 2.1: From Images to High-Fidelity 3D Assets with Production-Ready PBR Material. *arXiv preprint arXiv:2506.15442*.
- Kurai, R.; Hiraki, T.; Hiroi, Y.; Hirao, Y.; Perusquía-Hernández, M.; Uchiyama, H.; and Kiyokawa, K. 2025. MagicCraft: Natural Language-Driven Generation of Dynamic and Interactive 3D Objects for Commercial Meta-verse Platforms. *arXiv preprint arXiv:2504.21332*.
- Lai, Z.; Zhao, Y.; Liu, H.; Zhao, Z.; Lin, Q.; Shi, H.; Yang, X.; Yang, M.; Yang, S.; Feng, Y.; et al. 2025. Hunyuan3D 2.5: Towards High-Fidelity 3D Assets Generation with Ultimate Details. *arXiv preprint arXiv:2506.16504*.
- Li, P.; Ma, S.; Chen, J.; Liu, Y.; Zhang, C.; Xue, W.; Luo, W.; Sheffer, A.; Wang, W.; and Guo, Y. 2025a. CMD: Controllable Multiview Diffusion for 3D Editing and Progressive Generation. *arXiv preprint arXiv:2505.07003*.
- Li, W.; Liu, J.; Yan, H.; Chen, R.; Liang, Y.; Chen, X.; Tan, P.; and Long, X. 2025b. CraftsMan3D: High-fidelity Mesh Generation with 3D Native Diffusion and Interactive Geometry Refiner. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 5307–5317.
- Li, W.; Zhang, X.; Sun, Z.; Qi, D.; Li, H.; Cheng, W.; Cai, W.; Wu, S.; Liu, J.; Wang, Z.; et al. 2025c. Step1x-3d: Towards high-fidelity and controllable generation of textured 3d assets. *arXiv preprint arXiv:2505.07747*.
- Lim, B.; and Lee, S.-W. 2024. TIFu: Tri-Directional Implicit Function for High-Fidelity 3D Character Reconstruction. In *International Conference on Pattern Recognition and Artificial Intelligence*, 151–165. Springer.
- Lionar, S.; Liang, J.; and Lee, G. H. 2025. Treemeshgpt: Artistic mesh generation with autoregressive tree sequencing. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 26608–26617.
- Liu, J.; Xu, J.; Guo, S.; Li, J.; Guo, J.; Yu, J.; Weng, H.; Lei, B.; Yang, X.; Chen, Z.; et al. 2025. Mesh-RFT: Enhancing Mesh Generation via Fine-grained Reinforcement Fine-Tuning. *arXiv preprint arXiv:2505.16761*.
- Liu, M.; Zeng, C.; Wei, X.; Shi, R.; Chen, L.; Xu, C.; Zhang, M.; Wang, Z.; Zhang, X.; Liu, I.; et al. 2024. Meshformer: High-quality mesh generation with 3d-guided reconstruction model. *Advances in Neural Information Processing Systems*, 37: 59314–59341.
- Melnik, A.; Alt, B.; Nguyen, G.; Wilkowski, A.; Wu, Q.; Harms, S.; Rhodin, H.; Savva, M.; Beetz, M.; et al. 2025. Digital twin generation from visual data: A survey. *arXiv preprint arXiv:2504.13159*.
- Pan, Y.; Liu, C.; Xu, S.; Tan, S.; and Yang, J. 2025. Vasa-rig: Audio-driven 3d facial animation with ‘live’ mood dynamics in virtual reality. *IEEE Transactions on Visualization and Computer Graphics*.
- Ravi, N.; Reizenstein, J.; Novotny, D.; Gordon, T.; Lo, W.-Y.; Johnson, J.; and Gkioxari, G. 2020. Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:2007.08501*.
- Shah, J.; Bikshandi, G.; Zhang, Y.; Thakkar, V.; Ramani, P.; and Dao, T. 2024. Flashattention-3: Fast and accurate attention with asynchrony and low-precision. *Advances in Neural Information Processing Systems*, 37: 68658–68685.
- Siddiqui, Y.; Alliegro, A.; Artemov, A.; Tommasi, T.; Sirigatti, D.; Rosov, V.; Dai, A.; and Nießner, M. 2024. Meshgpt: Generating triangle meshes with decoder-only transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 19615–19625.

- Sinha, S.; Khan, M. S.; Usama, M.; Sam, S.; Stricker, D.; Ali, S. A.; and Afzal, M. Z. 2025. MARVEL-40M+: Multi-Level Visual Elaboration for High-Fidelity Text-to-3D Content Creation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 8105–8116.
- Tang, J.; Chen, Z.; Chen, X.; Wang, T.; Zeng, G.; and Liu, Z. 2024a. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. In *European Conference on Computer Vision*, 1–18. Springer.
- Tang, J.; Li, Z.; Hao, Z.; Liu, X.; Zeng, G.; Liu, M.-Y.; and Zhang, Q. 2024b. Edgerunner: Auto-regressive auto-encoder for artistic mesh generation. *arXiv preprint arXiv:2409.18114*.
- Wang, B.; Lan, J.; and Li, F. 2024. MSG-Voxel-GAN: multi-scale gradient voxel GAN for 3D object generation. *Multi-media Tools and Applications*, 83(40): 88505–88522.
- Weng, H.; Wang, Y.; Zhang, T.; Chen, C.; and Zhu, J. 2024. Pivotmesh: Generic 3d mesh generation via pivot vertices guidance. *arXiv preprint arXiv:2405.16890*.
- Weng, H.; Zhao, Z.; Lei, B.; Yang, X.; Liu, J.; Lai, Z.; Chen, Z.; Liu, Y.; Jiang, J.; Guo, C.; et al. 2025. Scaling mesh generation via compressive tokenization. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 11093–11103.
- Wu, C.; Ji, J.; Wang, H.; Ma, Y.; Huang, Y.; Luo, G.; Fei, H.; Sun, X.; Ji, R.; et al. 2024a. Rg-san: Rule-guided spatial awareness network for end-to-end 3d referring expression segmentation. *Advances in Neural Information Processing Systems*, 37: 110972–110999.
- Wu, K.; Liu, F.; Cai, Z.; Yan, R.; Wang, H.; Hu, Y.; Duan, Y.; and Ma, K. 2024b. Unique3d: High-quality and efficient 3d mesh generation from a single image. *Advances in Neural Information Processing Systems*, 37: 125116–125141.
- Wu, S.; Lin, Y.; Zhang, F.; Zeng, Y.; Xu, J.; Torr, P.; Cao, X.; and Yao, Y. 2024c. Direct3d: Scalable image-to-3d generation via 3d latent diffusion transformer. *Advances in Neural Information Processing Systems*, 37: 121859–121881.
- Wu, S.; Lin, Y.; Zhang, F.; Zeng, Y.; Yang, Y.; Bao, Y.; Qian, J.; Zhu, S.; Cao, X.; Torr, P.; et al. 2025. Direct3D-S2: Gigascale 3D Generation Made Easy with Spatial Sparse Attention. *arXiv preprint arXiv:2505.17412*.
- Xiang, J.; Lv, Z.; Xu, S.; Deng, Y.; Wang, R.; Zhang, B.; Chen, D.; Tong, X.; and Yang, J. 2025. Structured 3d latents for scalable and versatile 3d generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 21469–21480.
- Xu, J.; Cheng, W.; Gao, Y.; Wang, X.; Gao, S.; and Shan, Y. 2024. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. *arXiv preprint arXiv:2404.07191*.
- Zhang, B.; Tang, J.; Niessner, M.; and Wonka, P. 2023. 3dshape2vecset: A 3d shape representation for neural fields and generative diffusion models. *ACM Transactions On Graphics (TOG)*, 42(4): 1–16.
- Zhang, J.; Xiong, F.; and Xu, M. 2024. 3D representation in 512-Byte: Variational tokenizer is the key for autoregressive 3D generation. *arXiv preprint arXiv:2412.02202*.
- Zhang, L.; Wang, Z.; Zhang, Q.; Qiu, Q.; Pang, A.; Jiang, H.; Yang, W.; Xu, L.; and Yu, J. 2024. Clay: A controllable large-scale generative model for creating high-quality 3d assets. *ACM Transactions on Graphics (TOG)*, 43(4): 1–20.
- Zhao, R.; Ye, J.; Wang, Z.; Liu, G.; Chen, Y.; Wang, Y.; and Zhu, J. 2025a. Deepmesh: Auto-regressive artist-mesh creation with reinforcement learning. *arXiv preprint arXiv:2503.15265*.
- Zhao, Z.; Lai, Z.; Lin, Q.; Zhao, Y.; Liu, H.; Yang, S.; Feng, Y.; Yang, M.; Zhang, S.; Yang, X.; et al. 2025b. Hunyuan3d 2.0: Scaling diffusion models for high resolution textured 3d assets generation. *arXiv preprint arXiv:2501.12202*.
- Zhou, Z.; Zhou, S.; Lv, Z.; Zou, M.; Tang, Y.; and Liang, J. 2024. A simple baseline for efficient hand mesh reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1367–1376.