


TreeMeshGPT: Artistic Mesh Generation with Autoregressive Tree Sequencing

Stefan Lionar^{1,2,3}, Jiabin Liang^{1,2,3}, Gim Hee Lee³

¹Sea AI Lab

²Garena

³National University of Singapore



Motivation


Generative 3D models can produce high-quality 3D assets




meshy.ai



lumalabs.ai/genie



hyper3d.ai




Over-tessellated
meshes

Not suitable for
downstream
applications

Motivation


Artist-created mesh topology




Suitable for downstream applications:

- Compact vertices and faces:
 - Optimize file size
- Good edge flow for clean deformation

Bad edge flow




Good edge flow



Source: <https://www.youtube.com/watch?v=Lip59doQQRk>

- Simplify UV mapping




Source: Autodesk Maya 2025

Previous methods: Autoregressive Transformer Decoder

**MeshGPT
(2023)**



**MeshAnything
(2024)**




Point Cloud

Different tokenizations:

- Naive: MeshXL
- VQ-VAE: MeshGPT, MeshAnything, PivotMesh
- Adjacent Mesh Tokenization (AMT): MeshAnythingV2
- EdgeRunner

Previous methods: Challenges

- Long sequences
 - **Naive tokenization:** 1 triangular face \rightarrow 9 tokens (100% compression rate)



- **EdgeRunner & AMT:** Compression rate of $\pm 46\%$

- Incomplete geometry and flipped normals


Original mesh



MeshAnything



MeshAnythingV2



Flipped normals cause shading artifacts

With flipped normals



Fixed flipped normals




Source: <https://www.youtube.com/shorts/GPj1nTKWrik>

TreeMeshGPT: Autoregressive Tree Sequencing


Sequence order is retrieved from growing tree structures

Mesh subset




(v_2, v_1)
 (v_0, v_2)

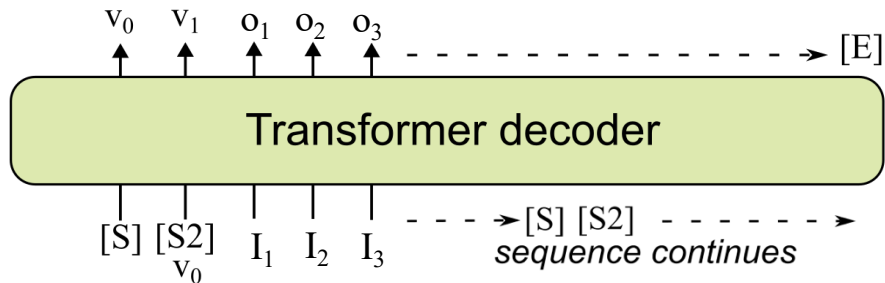
Dynamic stack (DFS traversal)



Tree structure

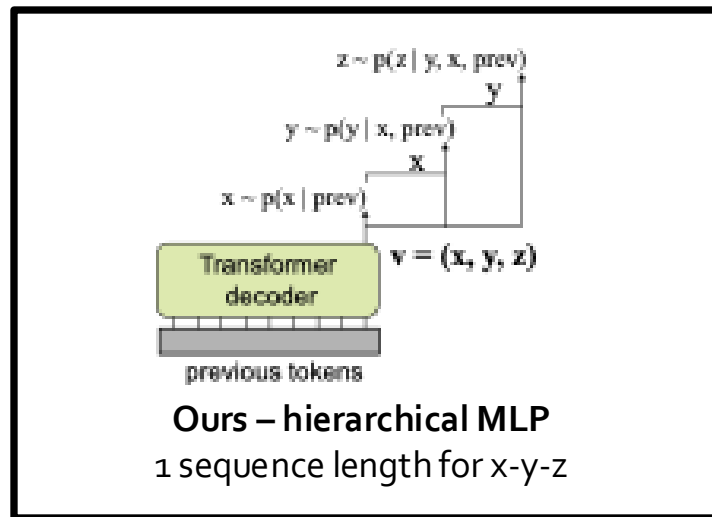
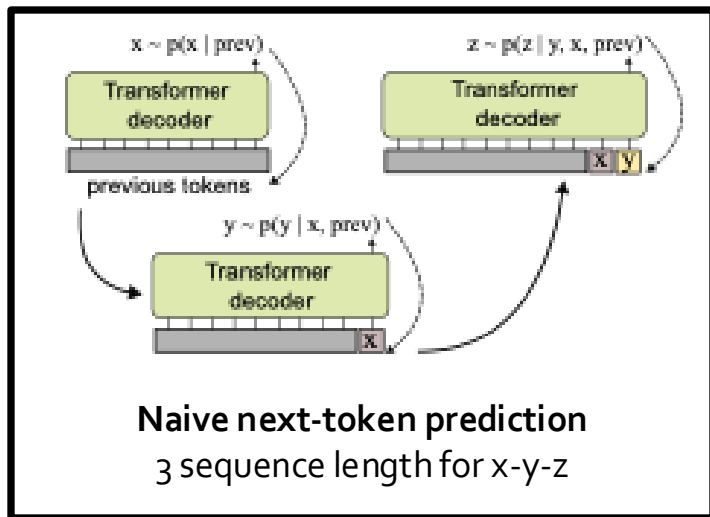


Autoregressive sequence



TreeMeshGPT: Autoregressive Tree Sequencing

Vertex Prediction




Advantages:


- 1) 1 face adds two tree nodes \rightarrow 22% compression rate
- 2) Local extension provides effective inductive bias
- 3) Strong normal constraint minimizing flipped normals

Results

Objaverse – Artist-created mesh dataset



GSO – Real-world scan dataset Best out of 5 trials



Model	CD↓	NC↑	NC ↑
MeshAnything	0.0105	0.453	0.869
MeshAnythingV2	0.0116	0.3269	0.865
Ours	0.0077	0.842	0.897

CD = Chamfer Distance

NC = Normal Consistency

|NC| = Normal Consistency ignoring flipping direction

Results

Text prompt

Luma Genie

Ours

Conditioned on point clouds sampled from dense meshes


"knight helmet with horns"



"anime cat with wings"



"cyberpunk car"



Results: 9-bit discretization





Thank you!



<https://github.com/sail-sg/TreeMeshGPT>