


EVA01: Unified Native 3D Understanding and Generation via Mixture-of-Transformers

SeeleAI Team

SEELE <https://www.seeles.ai/>
 <https://www.seeles.ai/research/pages/EVA01>

Abstract

This paper addresses the challenge of integrating 3D mesh as a native modality within Multimodal Large Language Models (MLLMs). Diffusion-based large reconstruction models decouple semantic understanding from geometric reasoning, operating as stateless reconstructors conditioned on dense 2D pixel priors. Recent MLLM-based methods treat the 3D modality as an external output rather than a native component of the multimodal sequence, making incremental adaptations without systematic analysis of how geometric manifolds align with MLLM feature spaces. We introduce EVA01, a unified framework that extends the modality boundary of MLLMs to natively incorporate 3D mesh understanding, generation, and context-aware editing. Built upon a Mixture-of-Transformers (MoT) architecture, EVA01 decouples the model into a pre-trained Understanding Expert (E_{und}) and a structurally mirrored Generation Expert (E_{gen}), coupled through shared global self-attention with hard modality routing. This design aligns the semantic latent space of the MLLM backbone with the geometric manifold, enabling direct transfer of multimodal priors without intermediate 2D representations. Results show that EVA01 achieves state-of-the-art native text-to-3D generation fidelity and unlocks robust long-context multi-turn geometric editing with identity preservation—a capability fundamentally inaccessible to stateless reconstruction pipelines. Our findings further offer architectural insights for integrating 2D foundation models with 3D tasks, informing the design of 3D-native multimodal systems.

Understanding

Native multi-turn reasoning over textured 3D meshes



Multi-turn Mesh QA

identity counting material local parts

User: What is this?

Assistant: This is a **robotic crab-like machine** with articulated legs and a central body.

User: How many legs does it have?

Assistant: The object has **six legs**.

User: What is the color of the legs?

Assistant: The legs are **red** with **white accents**.

User: What is on the back of the crab?

Assistant: The back features a **circular, gear-like structure**.

Context-Aware Generation & Editing




Figure 1. Native Multi-Turn 3D Understanding, Generation, and Editing. EVA01 unifies mesh-native 3D understanding, text-conditioned generation, and context-aware editing within a continuous context. Left: multi-turn question answering over a textured mesh. Right: text-to-3D generation followed by sequential edits across three assets. Each trajectory applies localized structural edits—revealing a face, adding weapons, unfolding wings, replacing tracks with wheels, attaching a robotic arm—while preserving object identity across turns. All edits are generated without explicit masks, conditioned on the full interaction history.

1 Introduction

Diffusion-based large reconstruction models have driven recent progress in 3D content creation by leveraging dense pixel-level features from multi-view images to reconstruct 3D geometry. These methods—spanning Score Distillation Sampling (SDS) Poole et al. (2023); Wang et al. (2023), cascaded multi-view reconstruction Xu et al. (2024a); Team (2025b), and feed-forward large reconstruction models Xiang et al. (2024)—exploit the spatial priors encoded in 2D foundation models Rombach et al. (2022); Saharia et al. (2022); Peebles and Xie (2023) to achieve high visual fidelity. However, they decouple semantic understanding from geometric reasoning: semantic interpretation is delegated to a frozen image encoder, while geometric construction is treated as a downstream lifting operation. This renders them geometry reconstructors conditioned on dense pixel priors rather than generative models that reason over 3D structure. Consequently, they operate in a stateless manner—every edit requires full re-generation, with no mechanism to preserve geometric identity across sequential modifications, limiting their utility for iterative 3D design.

Recent works have explored incorporating Multimodal Large Language Models (MLLMs) into 3D generation to bridge this semantic-geometric gap Ye et al. (2025b, 2026); Huang and Xu (2026); Chen et al. (2025b); Huang et al. (2026); Chen et al. (2026). While these approaches advance language-driven 3D generation, they treat the MLLM primarily as a semantic feature extractor or conditioning module without systematic representation-level analysis of the relationship between MLLM feature spaces and the 3D geometric manifold. Existing methods have not yet demonstrated a post-training pipeline—with modality alignment, progressive curriculum learning, and expert decoupling—that follows the scaling paradigm established by unified image understanding and generation models Deng et al. (2025); Xie et al. (2024); Wu et al. (2024). This limits their capacity to achieve robust semantic-geometric alignment and to scale with model capacity and data volume.

Native image generation models such as Nano-Banana and GPT-Image have recently demonstrated that treating images as first-class tokens within a unified sequence unlocks consistent, multi-turn understanding and editing inaccessible to prior approaches. Can the same transition be realized for 3D? This requires integrating 3D mesh as a first-class modality within the MLLM sequence stream, which introduces a representational challenge beyond simple multimodal feature concatenation. A unified sequence stream must accommodate three modalities with structurally distinct properties: text encodes abstract semantics with no spatial inductive bias; images capture dense pixel-level spatial correlations that implicitly encode projective geometry Oquab et al. (2023); Darcet et al. (2023); Siméoni et al. (2025); and 3D meshes impose strict topological constraints—manifoldness, genus, and local connectivity—that must be respected at every sequence position. The core question is therefore not merely how to generate 3D shapes from language, but how to align these three heterogeneous modalities within a unified sequence representation—where text and images follow autoregressive modeling while mesh geometry is generated via flow matching—such that semantic intent, visual grounding, and geometric validity are jointly preserved. This alignment problem is especially acute under the scarcity of large-scale, high-quality text–3D paired data Deitke et al. (2022, 2023); Zhang et al. (2025); Chang et al. (2015), demanding training strategies that efficiently transfer multimodal priors to the 3D domain.

We introduce **EVA01**, a context-aware unified MLLM that natively integrates 3D mesh understanding, generation, and multi-turn editing within a single Mixture-of-Transformers (MoT) architecture Liang et al. (2025). EVA01 decouples the model into two complementary experts: a pre-trained *Understanding Expert* (E_{und}) that serves as a stable semantic anchor preserving the multimodal priors of the MLLM backbone, and a structurally mirrored *Generation Expert* (E_{gen}) dedicated to geometry synthesis. Through shared global self-attention with hard modality routing, E_{gen} explicitly queries semantic representations from E_{und} , enabling cross-modal knowledge transfer while maintaining optimization independence. To address long-context 3D editing—where strict topological constraints must be preserved across sequential modifications—we construct a large-scale interleaved text–image–mesh dataset and formulate 3D generation as a **multimodal sequence modeling** task, where each geometric state is predicted conditioned on both the current instruction and the full historical context. This stateful formulation enables **context-aware, identity-preserving** multi-turn 3D editing inaccessible to stateless reconstruction pipelines.

Our contributions are threefold:

- **Unified MoT-based 3D MLLM:** To our knowledge, the first native 3D MLLM to combine mesh understanding, generation, and context-aware multi-turn editing within a single Mixture-of-Transformers architecture, integrating mesh as a first-class modality via a structurally mirrored generation expert and shared global self-attention for cross-modal knowledge transfer.
- **Curriculum-Based Semantic Alignment:** A multi-stage post-training strategy that bridges the mismatch between semantic representations and geometric structures via progressive modality alignment, using inter-

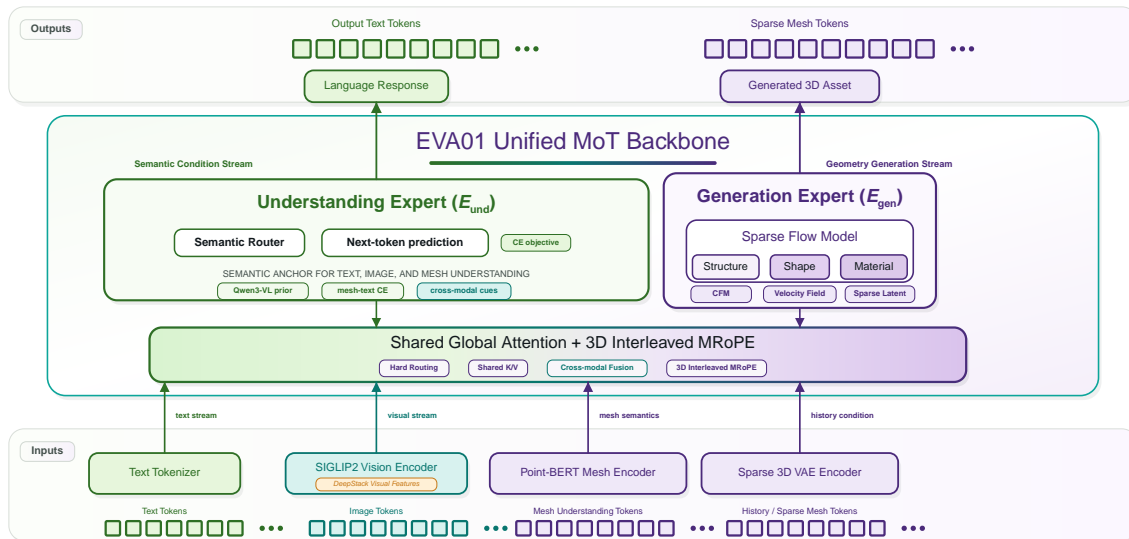


Figure 2. The Architecture of EVA01. EVA01 organizes tokenized text, image, and mesh inputs within a unified **Mixture-of-Transformers** backbone. The **Understanding Expert** (E_{und}) preserves the semantic priors of the pre-trained MLLM through Qwen tokenization, SigLIP2 visual encoding, Point-BERT mesh encoding, and DeepStack visual features. The **Generation Expert** (E_{gen}) consumes history-conditioned **Sparse Mesh Tokens** and predicts the structured 3D latent through **Structure**, **Shape**, and **Material** stages under Conditional Flow Matching. A shared fusion card, **Shared Global Attention + 3D Interleaved MRoPE**, enables cross-modal routing from semantic conditions to geometry generation while preserving spatial correspondence for context-aware editing.

leaved text–image–mesh sequences and modality dropout to establish robust cross-modal correspondences.

- **Stateful Editing Paradigm:** A stateful generation formulation that models 3D editing as conditional sequence modeling, achieving identity-preserving geometric modifications across multi-turn interactions and providing insights into the training dynamics of 3D-native MLLMs.

2 Related Work

2.1 3D Generative Models and Latent Representations

Generating high-fidelity 3D geometry at scale remains challenging due to data sparsity, cubic volumetric complexity, and strict topological requirements Kazhdan et al. (2006); Lorensen and Cline (1987). Optimization-based distillation (e.g., SDS) and cascaded multi-view reconstruction leverage 2D diffusion priors for 3D synthesis Poole et al. (2023); Wang et al. (2023); Xu et al. (2024a), but treat 3D as a reconstruction task conditioned on dense pixel priors rather than native generation, often producing view inconsistencies and decoupled semantic reasoning Team (2024, 2025b). Native 3D diffusion methods improve scalability by operating within learned latent manifolds, compressing structured signals via 3D VAEs for transformer-based generation Xiang et al. (2024); Team (2025b); Wu et al. (2025b); Jia et al. (2025). However, existing latent representations face a trade-off: global token sets oversmooth high-frequency details Zhang et al. (2023), while sparse localized tokens incur substantial computational cost Lai et al. (2025). Disjoint geometry–appearance modeling further causes semantic misalignment and fragile surfaces during remeshing, motivating unified, structurally-aware representations for coherent end-to-end generation Jia et al. (2025); Wu et al. (2025b); Lai et al. (2025).

2.2 Unified Multimodal Large Language Models

Multimodal Large Language Models (MLLMs) have evolved from LLMs augmented with modality-specific encoders to unified architectures that integrate generation within the language backbone, enabling fine-grained control and compositional reasoning across long contexts Deng et al. (2025); Dong et al. (2024); Xie et al. (2024); Wu et al. (2024); Ma et al. (2024); Chen et al. (2025a). Extending this paradigm to 3D, however, is constrained by geometric topology and the need for spatial consistency. Unlike 2D media, 3D assets require strict identity preservation across sequential modifications; cascaded text-to-image-to-3D pipelines decouple

semantic planning from geometric construction, leading to identity drift and topological discontinuities. Effective unification further demands latent representations expressive enough to capture high-frequency detail yet compatible with autoregressive prediction. Building a practical 3D-native multimodal model therefore requires synergizing MLLM semantic priors with scalable, topologically valid latent representations and production-ready data pipelines, enabling direct, context-aware, identity-preserving editing aligned with linguistic intent Wu et al. (2025b); Jia et al. (2025); Lai et al. (2025); Han et al. (2024).

2.3 Unified 3D Multimodal Large Models

ShapeLLM-Omni Ye et al. (2025b), built on Qwen2.5-VL-Instruct-7B, treats 3D as a first-class modality by expanding the LLM vocabulary with 8,192 learned 3D VQ-VAE tokens Xiang et al. (2024) and performing understanding and generation via fully autoregressive next-token prediction. Its single unified backbone processes text, image, and 3D tokens without modality-specific expert decoupling, causing pre-trained MLLM priors to degrade under the conflicting optimization demands of semantic reasoning and geometric synthesis; moreover, discrete VQ-VAE tokenization inherently limits geometric fidelity. Omni123 Ye et al. (2026) unifies text-to-2D and text-to-3D generation via dual-stream attention, Cube3D VQ-VAE tokenization, and dual text encoders (CLIP + Qwen3); it supports instruction-based editing and achieves strong results on Edit3D-Bench, yet each edit proceeds as an independent forward pass without persistent geometric identity across turns. CG-MLLM Huang and Xu (2026), built concurrently on Qwen3-VL, shares a similar MoT backbone with dedicated understanding and generation experts under hard modality routing, making it the closest architectural parallel to EVA01. The critical differences lie in representation and data: it relies on a VecSet-based 3D representation Team (2025a)—which our ablations (Sec. 4.7) show is substantially weaker than a structured sparse grid—performs mesh understanding on rendered 2D views rather than native mesh tokens, and lacks a multi-stage post-training curriculum with modality-specific optimization schedules. Sar3d Chen et al. (2025b) proposes next-scale autoregressive prediction over a multi-scale 3D VQ-VAE for fast generation (0.82s per object), and repurposes truncated token scales for captioning via a separate LLaMA; generation and understanding thus rely on distinct model components rather than a natively unified MLLM. UniMesh Huang et al. (2026) employs Bagel Deng et al. (2025) as a frozen MLLM conditioner—outputting FLUX VAE image latents bridged through a LoRA-fine-tuned Mesh Head to Hunyuan3D’s implicit shape decoder—and introduces Chain-of-Mesh for inference-time iterative editing. Know3D Chen et al. (2026) follows a three-stage pipeline: Qwen2.5-VL for semantic reasoning, Qwen-Image-Edit-2511 (20B MMDiT) for back-view generation, and frozen TRELIS.2 with parallel cross-attention injection of intermediate MMDiT hidden states Xiang et al. (2024). In both cases the MLLM serves as an external conditioner rather than a native reasoning backbone processing text, images, and 3D mesh tokens within a shared sequence space.

In summary, while these methods advance the integration of MLLMs with 3D content, none combine three capabilities that define EVA01: (1) a MoT architecture that decouples semantic understanding from geometric generation with shared cross-modal attention, (2) native 3D mesh tokens processed within the MLLM sequence stream via a structured grid-based latent representation and flow matching, and (3) stateful, long-context multi-turn 3D editing where each geometric update is conditioned on the full interaction history with explicit identity preservation.

3 Methodology

EVA01 is a unified mesh understanding and generation MLLM that natively processes text, images, and 3D geometry within a single Mixture-of-Transformers sequence stream. This section formalizes the sparse-voxel-based mesh tokenization, the MoT backbone with decoupled understanding and generation experts, and the conditional flow matching formulation that enables semantic-geometric alignment and context-aware generation.

3.1 Architecture: Unified 3D Multimodal Mixture-of-Transformers

EVA01 extends a pre-trained MLLM backbone (Qwen3-VL Bai et al. (2025)) to model 3D mesh geometry as a conditional flow matching problem within a unified sequence stream. Given a multimodal input sequence $\mathbf{X} = [\mathbf{x}_{\text{txt}}, \mathbf{x}_{\text{img}}, \mathbf{x}_{\text{mesh}}]$, where $\mathbf{x}_{(\cdot)}$ denotes the tokenized sequence of each modality, we learn the conditional probability distribution $p(\mathbf{x}_{\text{mesh}} \mid \mathbf{x}_{\text{txt}}, \mathbf{x}_{\text{img}}, \mathbf{c}_{\text{ctx}})$, with \mathbf{c}_{ctx} encoding the historical context in multi-turn interactions.

3D Mesh Tokenization. Following TRELIS.2 Xiang et al. (2025), we represent 3D assets as sparse voxel-based structures (O-Voxel) that jointly encode geometry and appearance. Each asset is defined as a collection

Table 1. Unified Block Attention Masking for Multi-Turn Editing. Visibility constraints within a packed sequence. Purple : Causal; Green : Bidirectional; Light gray : Masked. The staged 3D latent blocks consist of sparse structure (SS; dense latent), sparse shape (shape; sparse latent), and sparse material (material; sparse latent). The current generation conditions on clean historical geometry while noisy blocks are hidden from all later blocks.

Query ↓ Key →	Text 1	Image 1	Noise SS 1	Clean SS 1	Noise Shape 1	Clean Shape 1	Noise Material 1	Clean Material 1	Text 2	Noisy SS 2
Text 1	Causal	-	-	-	-	-	-	-	-	-
Image 1	Full	Causal	-	-	-	-	-	-	-	-
Noise SS 1	Full	Full	Bi-Dir	-	-	-	-	-	-	-
Clean SS 1	Full	Full	-	Bi-Dir	-	-	-	-	-	-
Noise Shape 1	Full	Full	-	Full	Bi-Dir	-	-	-	-	-
Clean Shape 1	Full	Full	-	Full	-	Bi-Dir	-	-	-	-
Noise Material 1	Full	Full	-	Full	-	Full	Bi-Dir	-	-	-
Clean Material 1	Full	Full	-	Full	-	Full	-	Bi-Dir	-	-
Text 2	Full	Full	-	Full	-	Full	-	Full	Causal	-
Noisy SS 2	Full	Full	-	Full	-	Full	-	Full	Full	Bi-Dir

of feature tuples anchored to a regular 3D grid of resolution N^3 :

$$\mathbf{f} = \{(\mathbf{f}_i^{\text{shape}}, \mathbf{f}_i^{\text{mat}}, \mathbf{p}_i)\}_{i=1}^L, \quad (1)$$

where $\mathbf{f}_i^{\text{shape}}$ encodes local geometric information (dual vertex position, edge intersection flags, and splitting weights), $\mathbf{f}_i^{\text{mat}}$ encodes PBR material parameters (base color, metallic, roughness, opacity), and $\mathbf{p}_i \in \{0, \dots, N-1\}^3$ is the coordinate of the i -th active voxel; inactive voxels are discarded. This representation supports direct bidirectional conversion to and from meshes via a Flexible Dual Grid formulation, avoiding the iterative decoding and field extraction of prior representations. A pre-trained VAE compresses the O-Voxel representation in Eq. 1 into compact sparse latent tokens $\mathbf{x}_{\text{mesh}} \in \mathbb{R}^{L \times C}$, where L is the number of active voxels.

EVA01 adopts a structured sparse grid representation in place of the VecSet paradigm Zhang et al. (2023) common in recent 3D generative frameworks. VecSet compresses point cloud data into order-invariant latent tokens via a Perceiver-like Jaegle et al. (2021) architecture; while compact, these unordered tokens lack explicit spatial coordinates, depriving attention mechanisms of the positional grounding required to establish stable geometric correspondences. Our sparse voxel representation anchors each token to a fixed coordinate \mathbf{p}_i , binding geometry to a regular spatial lattice—a structural regularity indispensable for long-context interleaved sequences, where sequence position and spatial topology must be tracked simultaneously. The sparse convolutional backbone further enables mesh sampling at substantially higher resolutions than VecSet methods, yielding reconstruction precision well beyond prior approaches. Ablation experiments (Sec. 4.7) confirm that explicit coordinate injection is a prerequisite for preventing geometric collapse in native generation.

Mixture-of-Transformers Backbone. To resolve the optimization dichotomy between semantic reasoning and geometric generation, we construct a multimodal MoT architecture Liang et al. (2025) derived from the Qwen3-VL Bai et al. (2025) backbone. We strictly decouple model parameters—spanning FFNs and attention projections—into two specialized sets: an Understanding Expert (E_{und}) and a Generation Expert (E_{gen}), as shown in Fig. 2.

Dual-Expert Routing. We employ a deterministic hard-routing strategy where computation for the i -th token is governed by its modality m_i (with $\text{txt}, \text{img} \mapsto E_{\text{und}}$ and $\text{mesh} \mapsto E_{\text{gen}}$). Formally, any linear transformation is computed via direct parameter indexing using modality-specific weight matrices $\mathbf{W}^{(m_i)}$:

$$\mathbf{h}'_i = \mathbf{h}_i \mathbf{W}^{(m_i)}. \quad (2)$$

This formulation segregates optimization: E_{und} remains a stable semantic anchor preserving the pre-trained MLLM priors, while E_{gen} undergoes geometric optimization without compromising understanding.

Understanding Expert Design. Following Janus Wu et al. (2025a), E_{und} decouples modality-specific encoding from the shared semantic backbone. Text is tokenized via the Qwen tokenizer, images are encoded through SigLIP2 Tschannen et al. (2025), and 3D mesh geometry is processed by a pre-trained Point-BERT Xu et al. (2024b) encoder. These modality-specific representations are projected into a shared semantic space derived from Qwen3-VL Bai et al. (2025), whose pre-aligned multimodal feature manifold exhibits 3D spatial understanding. This design preserves E_{und} as a semantic anchor without requiring modality-specific fine-tuning of the encoder stacks.

Generation Expert Design. E_{gen} implements a three-stage flow matching pipeline that mirrors the hierarchical structure of the sparse voxel representation. The **sparse structure** stage operates on a dense latent grid and

predicts the occupancy layout—which voxels are active—via a decoder with progressive upsampling, yielding the coordinate scaffolding for subsequent stages. Conditioned on this sparse layout, the **sparse geometry** stage generates shape features $\mathbf{f}^{\text{shape}}$ within active voxels using sparse convolutions and attention restricted to occupied coordinates. Finally, the **sparse material** stage synthesizes PBR material features \mathbf{f}^{mat} aligned to the generated geometry, similarly employing sparse operators. This ensures computational cost for the latter two stages scales with surface area rather than volume. At inference, the stages execute sequentially (structure \rightarrow geometry \rightarrow material); during training, they are parallelized via the unified attention mask (Table 1).

Shared Global Attention. To enable cross-modal reasoning despite parameter isolation, we implement shared global attention. For a token i with modality m_i , the output \mathbf{y}_i is computed as:

$$\mathbf{q}_i, \mathbf{k}_i, \mathbf{v}_i = \mathbf{h}_i \mathbf{W}_Q^{(m_i)}, \mathbf{h}_i \mathbf{W}_K^{(m_i)}, \mathbf{h}_i \mathbf{W}_V^{(m_i)} \quad (3)$$

$$\mathbf{y}_i = \text{Attn}(\mathbf{q}_i, \mathbf{K}, \mathbf{V}; \mathbf{M}) \mathbf{W}_O^{(m_i)} \quad (4)$$

where \mathbf{K}, \mathbf{V} aggregate keys and values from all sequence tokens j using their respective expert weights $\mathbf{W}_{\{K,V\}}^{(m_j)}$. The unified mask \mathbf{M} (Table 1) controls information flow, enabling the generation expert to query semantic priors from the understanding expert.

3D MRoPE. Standard 1D positional encodings flatten geometric structures, disrupting volumetric topology. Following Qwen3-VL Bai et al. (2025), we propose a 3D Interleaved MRoPE strategy that repurposes the original (T, W, H) rotary embeddings for sparse grid coordinates (x, y, z) in the generation branch. Rotary frequencies for the $X, Y,$ and Z axes are interleaved across the feature vector to maximize cross-dimension interaction, decomposing the head dimension D into strided subspaces indexed by $\mathcal{I}_x, \mathcal{I}_y, \mathcal{I}_z$ (e.g., $\mathcal{I}_x = \{k \mid k \equiv 0 \pmod{3}\}$). For a mesh token at sparse voxel coordinate $\mathbf{p} = (x, y, z)$, the rotary embedding is applied as:

$$\text{RoPE}(\mathbf{x}, \mathbf{p}) = \text{Interleave}(\mathcal{R}_x(\mathbf{x}_{\mathcal{I}_x}), \mathcal{R}_y(\mathbf{x}_{\mathcal{I}_y}), \mathcal{R}_z(\mathbf{x}_{\mathcal{I}_z})), \quad (5)$$

where \mathcal{R}_ϕ denotes rotation by spatial frequency ϕ , and $\mathbf{x}_{\mathcal{I}}$ represents the feature slice corresponding to subspace indices. This distributes spatial information uniformly across the sparse grid, injecting Euclidean inductive biases that prevent geometric drift during long-context editing.

3.2 Data: From Static Assets to Contextual Editing

To resolve the scarcity of paired 3D-text data and enable context-aware editing, we curate a hierarchical dataset in two phases: first assembling a large-scale corpus of high-fidelity static 3D assets, then synthesizing multi-turn interleaved editing trajectories from this foundation.

Phase 1: High-Quality Static 3D Asset Curation. We aggregate approximately 1.2M raw 3D assets from Objaverse-XL Deitke et al. (2023), TexVerse Zhang et al. (2025), and licensed internal repositories, spanning household objects, vehicles, characters, and mechanical assemblies. A rigorous standardization pipeline processes the full corpus while identifying a premium subset of 400K assets through stricter quality criteria: 1) *Geometric Standardization:* We canonicalize mesh orientations to a consistent world frame, verify UV coordinate integrity, validate texture resolution, and repair non-manifold edges and degenerate faces. Assets with irrecoverable geometric defects are discarded. 2) *Aesthetic and Structural Filtering:* A learned aesthetic scoring model trained on human preference annotations assigns quality ratings based on geometric complexity, texture realism, and visual appeal. The top 400K assets form a high-quality subset reserved for later training stages, while the full 1.2M corpus is retained for foundational pretraining. Topological analysis further tags objects with distinct structural properties—rigid sub-meshes amenable to part-based editing, articulated components with defined joint hierarchies, and assets with skeletal animation data—enabling targeted downstream editing tasks. 3) *Multi-View Dense Captioning:* Each asset is rendered from 8–12 uniformly sampled camera viewpoints under varying HDR environment lighting. A VLM processes these renderings to produce captions at three granularities: a short category-level summary, a medium-length description of geometric structure and material properties, and a detailed paragraph-level caption covering fine-grained shape details, surface texture, and functional affordances. These multi-granularity captions serve distinct training objectives across curriculum stages. The resulting dataset is formalized as triplets $\mathcal{D}_{\text{static}} = \{(\mathbf{t}_i, \{\mathbf{I}_{i,m}\}_{m=1}^M, \mathbf{x}_{\text{mesh},i})\}$, where \mathbf{t}_i is the caption, $\{\mathbf{I}_{i,m}\}$ the multi-view renders, and $\mathbf{x}_{\text{mesh},i}$ denotes the sparse voxel latent tokens comprising sparse structure, geometry, and material latent sets.

Phase 2: Interleaved Editing Sequences. We synthesize large-scale multi-turn editing sequences via two complementary pipelines, generating approximately 3M procedural and 300K semantic editing trajectories from the curated static asset pool. *Procedural Editing:* We algorithmically generate editing sequences using a composite operation set $\mathcal{O}_{\text{proc}} = \mathcal{O}_{\text{rigid}} \cup \mathcal{O}_{\text{anim}} \cup \mathcal{O}_{\text{topo}}$. $\mathcal{O}_{\text{rigid}}$ spans 6-DoF affine transformations (translation,

rotation, non-uniform scaling) applied to individual sub-mesh components, enabling part rearrangement, structural resizing, and component duplication. $\mathcal{O}_{\text{topo}}$ introduces topological perturbations including mesh boolean operations (union, difference, intersection) and localized mesh deformation fields, modeling constructive and destructive editing actions (e.g., adding a handle, carving a cavity). $\mathcal{O}_{\text{anim}}$ derives editing pairs from continuous animation sequences—skeletal deformations, mechanical articulations, and blend-shape morphs—by sampling frame pairs $(M_t, M_{t+\Delta t})$ with variable temporal stride Δt . Larger strides encourage learning long-range kinematic chains, while smaller strides provide fine-grained deformation supervision. *Semantic Editing*: An LLM generates diverse editing instructions \mathbf{t}_{inst} for each source mesh M_0 , covering attribute changes (material, color, texture), structural modifications (shape deformation, part replacement), and stylistic transformations. A conditioned image editing model modifies a representative view \mathbf{I}_0 to \mathbf{I}' following \mathbf{t}_{inst} , which is then lifted to 3D via a reconstruction pipeline with multi-view consistency enforcement. This pathway transfers semantic priors from 2D generative models into the 3D editing domain. The interleaved dataset is formalized as $\mathcal{D}_{\text{interleaved}} = \{\mathcal{S}^{(k)}\}$, where each sequence $\mathcal{S} = [(\mathbf{t}_0, \mathbf{I}_0), \mathbf{x}_{\text{mesh},0}, (\mathbf{t}_{\text{inst},1}, \mathbf{I}_{\text{ref},1}), \mathbf{x}_{\text{mesh},1}, \dots, (\mathbf{t}_{\text{inst},T}, \mathbf{I}_{\text{ref},T}), \mathbf{x}_{\text{mesh},T}]$ spans T editing turns, with each turn conditioning on the accumulated geometric and semantic context. Figure 3 summarizes this static-to-interleaved data construction pipeline.

3.3 Training: Multi-Stage Curriculum Learning

To bridge the misalignment between textual semantics and 3D topology, we employ a curriculum learning strategy that progressively aligns modalities and enables increasingly complex capabilities. The training objective combines two losses. For 3D generation, we formulate mesh synthesis as a Conditional Flow Matching (CFM) problem Lipman et al. (2023). Let $\mathbf{x}_1 \sim q(\mathbf{x}_{\text{mesh}})$ denote the clean sparse voxel latent and $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ standard Gaussian noise. Defining the probability path $\mathbf{x}_t = (1-t)\mathbf{x}_0 + t\mathbf{x}_1$, the flow matching loss is:

$$\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}_{t, \mathbf{x}_0, \mathbf{x}_1, \mathbf{c}} \left[\|\mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c}) - (\mathbf{x}_1 - \mathbf{x}_0)\|^2 \right], \quad (6)$$

where \mathbf{c} denotes the conditional context. For mesh understanding, we employ a standard autoregressive cross-entropy loss over text tokens conditioned on mesh features:

$$\mathcal{L}_{\text{CE}}(\theta) = - \sum_{i=1}^T \log p_\theta(t_i | t_{<i}, \mathbf{x}_{\text{mesh}}), \quad (7)$$

where $\{t_i\}$ are text caption tokens and \mathbf{x}_{mesh} is the sparse voxel latent of the corresponding 3D asset. Optimization proceeds through five stages, each introducing a distinct capability while preserving previously acquired knowledge.

Stage 1: Mesh Understanding Warm-up. We begin by establishing mesh understanding capability. Using mesh-text paired data from $\mathcal{D}_{\text{static}}$, we train only a lightweight MLP projector that maps Point-BERT Xu et al. (2024b) mesh features into the Qwen3-VL Bai et al. (2025) text embedding space. All other parameters—including E_{und} , E_{gen} , the mesh VAE, and the modality-specific encoders—remain frozen. This stage is optimized solely with \mathcal{L}_{CE} on the captioning objective, warming up the understanding pathway at minimal cost while preserving the pre-trained MLLM backbone.

Stage 2: Visual-Geometric Initialization. With a functioning understanding pathway in place, we establish generation capability by mapping dense visual features to the 3D manifold. Qwen3-VL’s DeepStack strategy Bai et al. (2025) propagates low-level visual features from early SigLIP2 Tschannen et al. (2025) layers directly into the LLM backbone, providing dense spatial and textural cues essential for accurate image-to-3D reconstruction. Ablation experiments (Sec. 4.7) confirm that removing this pathway significantly degrades geometric fidelity. We train E_{gen} with \mathcal{L}_{FM} using image-conditioned samples from $\mathcal{D}_{\text{static}}$, while E_{und} is jointly trained with mesh-text paired data under \mathcal{L}_{CE} . This dual-objective setup creates a synergistic effect: the mesh-text understanding signal sharpens E_{und} ’s geometric awareness, and through shared global self-attention, these refined semantic representations directly inform E_{gen} ’s image-to-3D generation, improving reconstruction fidelity beyond what image conditioning alone can achieve. A reduced learning rate on E_{und} preserves its pre-trained multimodal priors while enabling this cross-task transfer.

Stage 3: Semantic Modality Alignment. To bridge the textual and geometric manifolds, we introduce *Triple-Batch Sampling* combined with *Modality Dropout*. For each triplet in $\mathcal{D}_{\text{static}}$, we construct three independent training samples conditioned on text, images, and mesh-text pairs respectively. Dynamic image token dropout (p_{drop}) compels the network to reconstruct \mathbf{x}_{mesh} from textual cues alone, implicitly distilling the visual-geometric priors from Stage 2 into the text-conditioned generation process. Mesh-text pairing continues to reinforce understanding via \mathcal{L}_{CE} , while E_{und} and E_{gen} are jointly optimized with modality-specific learning rates to ensure output consistency across conditioning modalities.

Table 2. Training Recipe of EVA01. Multi-stage curriculum with differential optimization. E_{und} : Understanding Expert; E_{gen} : Generation Expert. MSE denotes the flow-matching regression loss. **Green highlight**: interleaved editing data.

	Stage 1 (Warm-up)	Stage 2 (Init)	Stage 3 (Align)	Stage 4 (Edit)	Stage 5 (HQ FT)
Hyperparameters					
Learning rate (E_{gen})	0.0	2e-4	1e-4	5e-5	5e-6
Learning rate (E_{und})	1e-4	1e-5	1e-5	5e-6	1e-6
LR scheduler	Constant	Constant	Cosine	Cosine	Constant
Optimizer	AdamW ($\beta_{1,2}=0.9, .95; \epsilon=10^{-8}; wd=0.05$)				
Training steps	20K	150K	100K	80K	30K
Sequence length	32-40K	56-64K	56-64K	80-96K	80-96K
Max context window	24K	30K	30K	60K	60K
Mesh resolution	–	512 ³	512 ³	512 ³	512 ³
Diffusion timestep shift	–	1.0	3.0	3.0	3.0
Loss weight (CE : MSE)	1:0	0.25:1	0.25:1	0.25:1	0:1
Data Sampling Ratio					
Image-Text Pair (I2T)	0.0	0.0	0.05	0.05	0.0
Image-Mesh Pair (I2M)	0.0	0.9	0.3	0.1	0.1
Text-Mesh Pair (T2M)	0.0	0.0	0.5	0.1	0.4
Text-Mesh Pair (M2T)	1.0	0.1	0.15	0.1	0.0
Procedural (Interleaved)	0.0	0.0	0.0	0.35	0.1
Semantic (Interleaved)	0.0	0.0	0.0	0.3	0.4

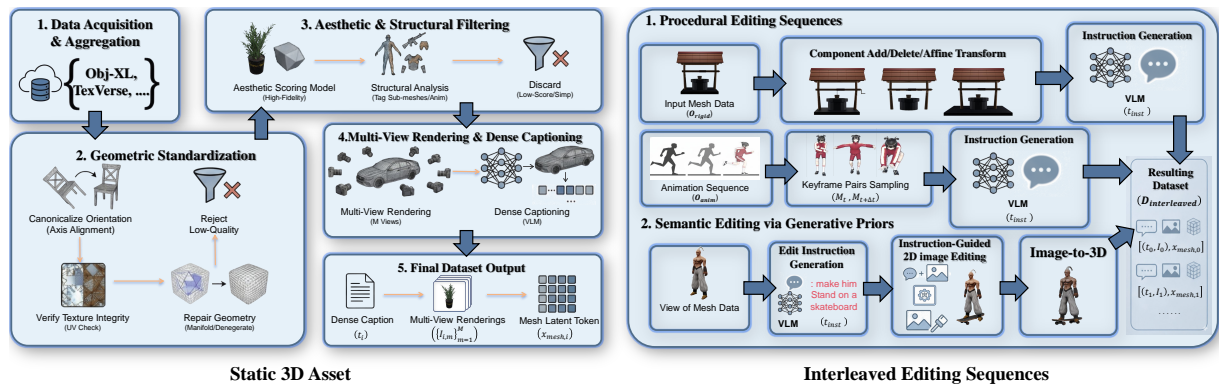


Figure 3. Data Curation Pipeline of EVA01. (Left) **Static 3D Asset Curation:** We standardize raw 3D assets through geometric canonicalization, aesthetic filtering, and multi-view dense captioning to construct high-quality text-image-mesh triplets. (Right) **Interleaved Editing Sequences:** To enable context-aware editing, we synthesize multi-turn sequences via two complementary pathways: *Procedural Editing* (top right) utilizing rigid transformations and animation keyframes for structural precision, and *Semantic Editing* (bottom right) leveraging 2D generative priors for open-ended stylistic modification.

Stage 4: Context-Aware Instruction Tuning. We enable sequential, stateful 3D editing using the interleaved dataset $\mathcal{D}_{interleaved}$. The generation of the k -th geometric state is conditioned on the full interaction history $\mathbf{c}_k = \{\mathbf{t}_{inst}, \mathbf{x}_{hist}\}$, with 3D Interleaved MRoPE encoding both global interaction timestamps and local spatial coordinates. The attention mask (Table 1) enforces that \mathbf{x}_k attends to the clean latent of \mathbf{x}_{k-1} , enabling differential editing operations—modifying topology based on instructions while preserving identity across turns—rather than stateless regeneration. This stage employs both \mathcal{L}_{FM} and \mathcal{L}_{CE} to maintain understanding quality alongside the emerging editing capability.

Stage 5: High-Quality Finetuning. The final stage refines generation fidelity on the curated 400K high-quality asset subset. Training continues with \mathcal{L}_{FM} at reduced learning rates, sharpening geometric details and surface quality. This stage is critical for elevating the upper bound of visual fidelity while preserving the context-aware editing behavior and mesh understanding capabilities acquired in prior stages.

4 Experiments

We evaluate EVA01 across four capability axes: single-turn generation (text-to-3D and image-to-3D), multi-turn context-aware editing, and mesh understanding. Beyond standard benchmarks, we report ablation studies dissecting key architectural and training components, analyze learned representations, and derive insights from training dynamics.

4.1 Experimental Protocol

Data. We follow the data curation pipeline described in Section 3.2, aggregating 1.2M raw assets (filtered to a 400K premium subset) for static training and synthesizing 3M procedural and 300K semantic editing sequences for multi-turn interactivity. Per-stage data sampling ratios are specified in Table 2.

For the understanding task, we build on the released PointLLM-V2 raw_v2 corpus, a point-native object-centric 3D-language dataset. We use its stage-structured split, namely Stage1-1M for point-language alignment and Stage2-700k for instruction tuning. This organization is well matched to our design, enabling disentanglement of faithful 3D semantic grounding from higher-level instruction-following behavior. In addition, we leverage our curated text-mesh pairs for caption training and employ the interleaved editing dataset to predict semantic editing instructions, making full use of our constructed data.

Implementation. EVA01 is built upon **Qwen3-VL-2B-Instruct** Bai et al. (2025) as the MLLM backbone. The Generation Expert (E_{gen}) is structurally mirrored from E_{und} and initialized from its pre-trained weights to accelerate convergence. The pre-trained mesh VAE and Point-BERT Xu et al. (2024b) encoder remain frozen throughout all training stages; the visual encoder SigLIP2 Tschannen et al. (2025) is trained. Training uses AdamW Loshchilov and Hutter (2019) ($\beta_1=0.9$, $\beta_2=0.95$, weight decay=0.05) with bfloat16 mixed precision and FlexAttention. A cosine learning rate schedule with linear warmup (1,000–5,000 steps, stage-dependent) is employed; per-stage learning rates and training steps are detailed in Table 2. Training uses sample packing with a batch size of 1—multiple samples concatenated into a single packed sequence—distributed across 32 NVIDIA H20 (96GB) GPUs via Fully Sharded Data Parallel (FSDP). Classifier-free guidance with a conditioning dropout rate of 0.1 is applied across all flow-matching stages. Total training wall-clock time is approximately two months.

Evaluation Benchmarks and Metrics. We evaluate across three axes. For **single-turn generation**, we use the standard **Toys4K** Stojanov et al. (2021) benchmark (3,218 high-quality 3D assets), reporting CLIP Score Hessel et al. (2021) for text-shape semantic alignment, and Fréchet Distance (FD) and Kernel Distance (KD) based on DINOv2 Oquab et al. (2023) features for geometric and visual fidelity. For **multi-turn editing**, we curate a 400-sample evaluation set from our interleaved editing corpus, selecting 200 procedural and 200 semantic editing sequences. Existing local-editing benchmarks such as Edit3D-Bench are useful references, but their limited scale (roughly 100 source objects), single-turn structure, and coarse edit-region coverage do not fully stress long-context mesh-native editing. For precise geometric edits, we provide manually verified masks to support mask-dependent baselines; for broad style and semantic edits, the masks indicate only approximate affected regions. EVA01 itself does not consume masks at inference time. We evaluate unedited-region consistency using Chamfer Distance (CD) and masked multi-view PSNR, and assess overall editing quality with CLIP, FD_{DINOv2} , and user-study preference (Pref%).

For **mesh understanding**, we evaluate on the PointLLM-200 Objaverse captioning benchmark Xu et al. (2024b), which contains 200 held-out Objaverse objects with official PointLLM reference captions from the human-annotated Cap3D split. We use the prompt “Caption this 3D model in detail.” under a prompt-only generation protocol: the model receives only the 3D input and the captioning prompt, while the reference caption is used solely for scoring. This avoids teacher-forcing or continuation-style leakage where the reference answer is included in the model context.

We report lexical, semantic, and judge-based captioning metrics. BLEU- n Papineni et al. (2002), ROUGE-L Lin (2004), and METEOR Banerjee and Lavie (2005) are computed with standard implementations and reported on a 0–100 scale. Semantic agreement beyond token overlap is measured via cosine similarity between reference and generated captions using Sentence-BERT Reimers and Gurevych (2019) and SimCSE Gao et al. (2021), scaled by 100. We further report two complementary GPT-based judge metrics. **GPT-ref** follows the PointLLM protocol Xu et al. (2024b): the judge is given the reference caption and the model prediction, estimates the fraction of reference aspects correctly or partially covered by the prediction, and returns a score in [0, 100]

Table 3. Quantitative comparisons on Toys4K. We report CLIP, FD_{DINOv2} , KD_{DINOv2} , and user-study preference (Pref). KD is reported $\times 100$. N/A denotes unsupported modalities.

Method	Text-to-3D				Image-to-3D			
	CLIP ↑	FD_{DINOv2} ↓	KD_{DINOv2} ↓	Pref ↑	CLIP ↑	FD_{DINOv2} ↓	KD_{DINOv2} ↓	Pref ↑
Text-to-3D and Image-to-3D								
Shap-E Jun and Nichol (2023)	25.12	495.23	49.85	0.0%	82.15	465.11	62.55	0.0%
3DTopia-XL Chen et al. (2025c)	25.85	425.12	35.23	0.0%	83.42	380.88	45.15	0.0%
TRELLIS Xiang et al. (2024)	<u>30.80</u>	<u>238.45</u>	<u>4.25</u>	<u>14.8%</u>	85.65	67.32	0.75	5.1%
ShapeLLM-Omni Ye et al. (2025b)	26.70	310.55	18.20	5.1%	84.50	145.32	12.45	0.0%
Michelangelo Zhao et al. (2023)	25.92	405.36	27.40	0.0%	83.20	360.45	42.20	0.0%
Text-to-3D only								
3DGen-R1 Tang et al. (2025a)	29.35	263.72	6.85	7.9%	N/A	N/A	N/A	N/A
GVGEN He et al. (2024)	25.55	438.90	30.65	1.9%	N/A	N/A	N/A	N/A
Image-to-3D only								
TRELLIS.2 Xiang et al. (2025)	N/A	N/A	N/A	N/A	89.34	56.82	0.49	41.7%
Hunyuan3D-2.1 Team (2025a)	N/A	N/A	N/A	N/A	86.95	64.80	0.66	10.2%
Direct3D-S2 Wu et al. (2026)	N/A	N/A	N/A	N/A	86.05	72.25	0.88	6.9%
Step1X-3D Li et al. (2025b)	N/A	N/A	N/A	N/A	86.60	66.10	0.72	7.9%
Hi3DGen Ye et al. (2025a)	N/A	N/A	N/A	N/A	86.48	65.70	0.70	6.0%
Ours	35.72	122.48	1.18	70.4%	<u>87.28</u>	<u>61.74</u>	<u>0.63</u>	<u>22.2%</u>

averaged over all valid responses. Since PointLLM-200 references are often brief single-sentence captions, we additionally report **GPT-img**, a render-grounded judge score. For GPT-img, GPT-5.5 is given four RGB renders of the object (front, right, back, and left) together with the candidate caption, but no reference caption, and scores how faithfully the caption describes the visible 3D object. All metrics are computed with the same scorer within each metric family, and higher is better.

Baselines. For single-turn generation, we compare against Shap-E Jun and Nichol (2023), 3DTopia-XL Chen et al. (2025c), TRELLIS Xiang et al. (2024), Michelangelo Zhao et al. (2023), 3DGen-R1 Tang et al. (2025a), GVGEN He et al. (2024), TRELLIS.2 Xiang et al. (2025), Hunyuan3D-2.1 Team (2025a), Direct3D-S2 Wu et al. (2026), Step1X-3D Li et al. (2025b), Hi3DGen Ye et al. (2025a), and ShapeLLM-Omni Ye et al. (2025b). For multi-turn editing, we compare against Instant3DiT Barda et al. (2025), TRELLIS Xiang et al. (2024), and VoxHammer Li et al. (2025a). For mesh understanding, we compare against PointLLM-7B/13B Xu et al. (2024b), ShapeLLM-7B/13B Qi et al. (2024), MiniGPT-3D Tang et al. (2024), LLaMA-Mesh Wang et al. (2024), GreenPLM Tang et al. (2025b), and ShapeLLM-Omni Ye et al. (2025b). All baselines are evaluated using their released checkpoints and official inference protocols.

4.2 Single-Turn Generation

Table 3 and Figure 4 evaluate single-turn generation under both text- and image-conditioned settings on Toys4K. These two settings stress complementary capabilities: text-to-3D requires mapping abstract language to plausible geometry without dense spatial evidence, whereas image-to-3D rewards faithful reconstruction

from pixel-aligned visual cues. EVA01 is designed to support both regimes within one mesh-native sequence model, rather than relying on separate task-specific pipelines.

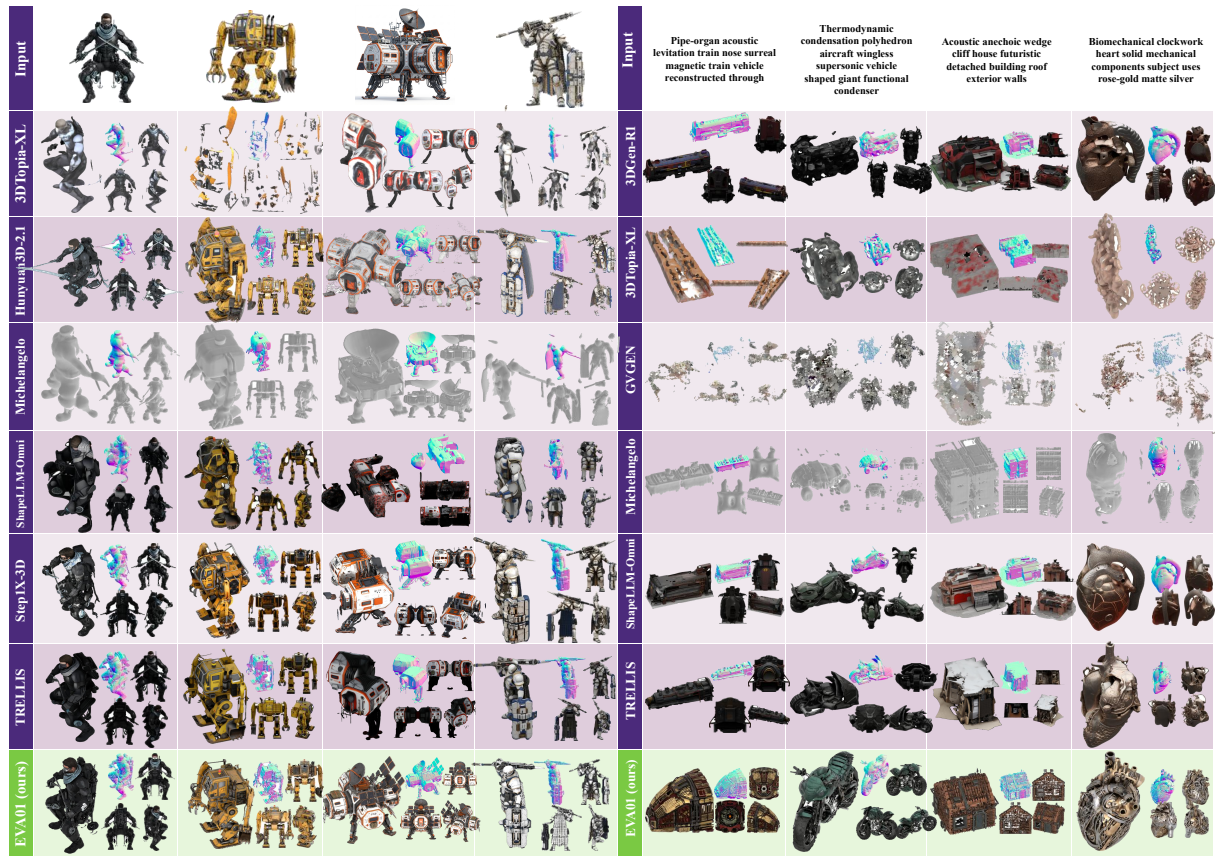


Figure 4. Qualitative Comparison with Baselines. We compare EVA01 with representative text-to-3D and image-to-3D baselines on Toys4K, including 3DTopia-XL, Hunyuan3D-2.1, Michelangelo, ShapeLLM-Omni, Step1X-3D, TRELIS, 3DGen-R1, and GVGEN. Across both image-conditioned cases (left) and text-conditioned cases (right), EVA01 better preserves object-level semantics, part structure, and material consistency, producing complete meshes with coherent geometry where prior methods often suffer from missing components, distorted topology, over-smoothed shapes, or fragmented surfaces.

Bridging the Semantic Gap in Text-to-3D. In text-to-3D, EVA01 achieves **35.72 CLIP**, **122.48 FD**, **1.18 KD**, and **70.4%** preference, outperforming the strongest baseline TRELIS (30.80 CLIP, 238.45 FD, 4.25 KD, 14.8% preference) and the text-specialized 3DGen-R1 (29.35 CLIP, 263.72 FD, 6.85 KD, 7.9% preference). This large margin indicates that dense-reconstruction backbones, although effective when visual evidence is available, remain inefficient at converting discrete language semantics into coherent 3D structure. EVA01 narrows this semantic-to-geometric gap by routing language-conditioned reasoning through the Understanding Expert (E_{und}), while the Generation Expert (E_{gen}) synthesizes sparse 3D latents through our staged alignment bridge and image-to-mesh warm-up curriculum. Compared with autoregressive mesh-token generation such as ShapeLLM-Omni (FD 310.55, KD 18.20), the flow-matching decoder further avoids severe quantization artifacts and yields smoother, more complete topology.

The right block of Figure 4 provides a more diagnostic view of this gap. The prompts contain unusual compositions and long attribute chains, such as a pipe-organ levitation train, a wingless supersonic vehicle, an anechoic wedge-cliff house, and a biomechanical clockwork heart. Baselines tend to satisfy only fragments of these descriptions: GVGEN often degenerates into sparse disconnected pieces, Michelangelo produces over-smoothed or incomplete gray geometry, and reconstruction-oriented methods such as 3DTopia-XL, ShapeLLM-Omni, and TRELIS frequently recover local parts while missing the intended global object category or material organization. EVA01 is not merely sharper at the surface level; it more consistently preserves the requested object identity, assembles semantically related parts into a coherent whole, and produces paired texture/normal outputs that remain structurally aligned. This behavior supports the role of our bridge training: language semantics are first grounded in the MLLM representation space and then transferred to the sparse 3D latent space, rather than being learned as a weak text condition attached to a reconstruction model.

High-Fidelity Image-to-3D. For image-to-3D, TRELIS.2 remains the strongest specialist model, achieving **89.34 CLIP**, **56.82 FD**, **0.49 KD**, and **41.7%** preference. EVA01 ranks second with 87.28 CLIP, 61.74 FD,

0.63 KD, and 22.2% preference, ahead of Hunyuan3D-2.1, Step1X-3D, Direct3D-S2, Hi3DGen, and the original TRELIS in the overall metric profile. This ranking is expected: image-to-3D is primarily a dense reconstruction problem, and TRELIS.2 is explicitly optimized to exploit pixel-aligned visual evidence from a single image. EVA01 instead uses the same MoT architecture and sparse latent interface for text generation, image generation, and downstream editing; its result therefore measures how much dense visual fidelity can be retained without abandoning a unified 3D-native MLLM formulation.

The left block of Figure 4 makes this trade-off visible. For character, robot, vehicle, and articulated object inputs, weak baselines either lose major structural parts, scatter geometry into disconnected fragments, or produce texture/normal predictions that no longer correspond to the same underlying shape. Michelangelo often recovers a coarse gray volume but weakens material and part separation; 3DTopia-XL and Hunyuan3D-2.1 can produce plausible local components but show instability in global assembly; ShapeLLM-Omni, Step1X-3D, and TRELIS preserve more recognizable structure yet still introduce missing limbs, duplicated subparts, or inconsistent normal maps in challenging cases. EVA01 generally maintains the input object’s category, pose, and large-scale topology while producing texture and normal renderings that remain aligned across views. The remaining gap to TRELIS.2 reflects the current limit of our SigLIP2/DeepStack visual path rather than a failure of the MoT formulation; we analyze this limitation and discuss concrete DINOv3-inspired improvements in Section 5.

4.3 Mesh Understanding

Table 4. Quantitative evaluation of mesh understanding on PointLLM-200. All models are evaluated on the PointLLM-200 Objaverse captioning benchmark with the same prompt: “Caption this 3D model in detail.” We report traditional lexical overlap metrics, embedding-based semantic metrics, and two GPT-5.5 judge scores. **GPT-ref** scores captions against the human reference text, while **GPT-img** scores captions directly against multi-view RGB renders without using the reference caption. All reported metrics are computed with the same scorer within each metric family, and higher is better.

Model	Traditional Metrics				Semantic Metrics		Judge Metrics	
	B-1 ↑	B-4 ↑	R-L ↑	METEOR ↑	SBERT ↑	SimCSE ↑	GPT-ref ↑	GPT-img ↑
GT caption	–	–	–	–	–	–	–	56.050
PointLLM-7B	7.905	0.669	10.290	13.536	47.894	48.812	<u>52.360</u>	51.495
PointLLM-13B	7.873	0.649	10.519	13.620	47.539	48.602	51.735	49.745
ShapeLLM-7B	11.491	1.222	14.035	13.316	37.693	38.047	24.560	30.075
ShapeLLM-13B	10.542	1.050	12.954	14.234	39.935	40.728	33.925	35.870
MiniGPT-3D	5.482	0.468	7.680	10.531	28.708	23.722	17.360	7.560
LLaMA-Mesh	5.119	0.441	7.439	9.904	26.326	25.095	10.740	7.085
GreenPLM	<u>17.527</u>	<u>2.303</u>	<u>20.632</u>	19.906	47.826	48.871	50.305	50.475
ShapeLLM-Omni	11.326	1.197	14.190	13.276	34.617	35.115	25.625	20.190
EVA01-Align	23.592	3.555	28.002	<u>19.007</u>	52.178	52.005	48.825	<u>56.770</u>
EVA01-Final	6.334	0.585	9.323	13.566	<u>50.554</u>	<u>50.878</u>	59.095	65.910

Table 4 reveals two complementary aspects of mesh understanding: reference-aligned captioning and open-ended semantic coverage. The model after the first three curriculum stages, denoted **EVA01-Align**, achieves the strongest reference-aligned captioning performance among all evaluated model predictions. It obtains the best BLEU-1, BLEU-4, ROUGE-L, Sentence-BERT, and SimCSE scores, and ranks second on METEOR. This confirms that the first three stages of our curriculum successfully establish a strong mesh-language alignment pathway. In particular, the mesh understanding warm-up, visual-geometric initialization, and semantic modality alignment stages jointly map 3D geometry into the language space in a way that matches the official PointLLM-200 reference caption distribution more effectively than prior 3D understanding baselines.

However, the table also highlights an important limitation of standard captioning metrics for evaluating mesh understanding. BLEU, ROUGE-L, and METEOR primarily measure lexical overlap with a single reference caption. As a result, they favor predictions that closely match the wording, length, and style of the ground-truth caption. This is useful for measuring reference-style caption imitation, but it does not fully capture open-ended 3D understanding: the same mesh can be correctly described with different object names, attributes, part descriptions, or levels of detail. A model may mention valid geometric or visual details that are visible in

the mesh but absent from the single reference caption, and such details can reduce n -gram precision despite being semantically correct. We therefore interpret these traditional metrics as measuring alignment to the PointLLM-200 caption style, rather than as complete measures of semantic mesh understanding.

This distinction explains the behavior of **EVA01-Final**. After the final instruction-oriented and high-quality finetuning stages, its lexical overlap scores decrease substantially, indicating that its output distribution shifts away from short, reference-like captions. At the same time, EVA01-Final retains the second-best Sentence-BERT and SimCSE scores among model predictions, and achieves the highest score under both GPT-ref and GPT-img. As an auxiliary sanity check, we also score the official reference caption itself with GPT-img. Both EVA01-Align and EVA01-Final score above this GT-caption row (56.77 and 65.91 vs. 56.05), where the judge compares each candidate caption directly against multi-view renders rather than against the reference sentence. This does not mean that the official reference captions are incorrect; rather, the PointLLM-200 references are brief human-annotated captions intended as single-reference evaluation targets, while GPT-img rewards detailed visual coverage of the rendered 3D object.

These results suggest that the later stages do not erase the mesh semantics learned by EVA01-Align. Instead, they repurpose this grounding for more instruction-following, fine-grained, and human-preferred descriptions. This behavior is consistent with prior findings in instruction tuning, where models trained to follow natural-language instructions become better aligned with user intent and human preference, even when their outputs diverge from narrow reference text patterns [Ouyang et al. \(2022\)](#); [Chung et al. \(2024\)](#). In our setting, the final model trades strict reference-caption overlap for broader semantic coverage, which is better captured by the GPT-based judges.

Among external baselines, GreenPLM [Tang et al. \(2025b\)](#) is the strongest reference-overlap competitor. It achieves the best baseline BLEU-1, BLEU-4, and ROUGE-L scores, and obtains the highest METEOR score overall. Under the GPT-based judges, however, PointLLM-7B is the strongest external baseline, while GreenPLM remains close. Nevertheless, EVA01-Align surpasses GreenPLM on most overlap and embedding-based metrics, while EVA01-Final exceeds all external baselines by a large margin under both reference-based and render-grounded GPT evaluation. The comparison with PointLLM also shows that simply scaling the language backbone is insufficient: PointLLM-7B and PointLLM-13B perform similarly across most metrics, suggesting that effective 3D-language alignment is more important than language-model size alone. Meanwhile, methods whose released inference interfaces are less matched to the PointLLM-200 point-cloud captioning protocol, including OBJ-text serialization in LLaMA-Mesh and the released 3D MLLM checkpoints of MiniGPT-3D and ShapeLLM variants, transfer less effectively under this unified prompt-only evaluation.

Overall, these results support the intended role of our curriculum. The first three stages produce a strong reference-aligned mesh captioner, demonstrating successful alignment between geometric input and language, but the qualitative examples in [Figures 5 and 6](#) show that EVA01-Align often stays at the level of object identity and coarse attributes. The final two stages adapt this understanding pathway to the broader unified setting, where the model must follow instructions and support generation/editing behavior rather than merely imitate short captions. As shown in [Figure 6](#), EVA01-Final produces richer descriptions of parts, materials, colors, and local structure, explaining its stronger GPT-img score. This gain comes with the usual instruction-tuning trade-off: the model may occasionally infer details not explicitly visible in the rendered input. The final model is therefore not optimized for maximum lexical overlap with a single reference, but for fine-grained, visually grounded semantic coverage.

4.4 Multi-Turn Context-Aware Editing

A defining capability of EVA01 is native, multi-turn 3D editing. Unlike cascaded pipelines that process each edit as an independent reconstruction problem, EVA01 models editing as a **stateful conditional sequence generation task**: the next mesh state is predicted from the current instruction and the accumulated text-image-mesh history. This formulation is essential for long-horizon editing, where the model must both apply a localized change and preserve the identity, topology, material layout, and previous edits of the same object.

Stateful Editing without Explicit Masks. EVA01 enforces the block attention mask in [Table 1](#), allowing the current noisy mesh tokens to attend to clean historical mesh states while preventing future leakage. Together with 3D Interleaved MRoPE, this gives each sparse latent token both a temporal role in the interaction and a spatial coordinate in the 3D grid. As a result, the model learns a differential geometric update conditioned on history, rather than regenerating the object from scratch. This is the main distinction from mask-conditioned baselines: Instant3DiT [Barda et al. \(2025\)](#) edits through multiview inpainting followed by reconstruction, TRELIS [Xiang et al. \(2024\)](#) uses a structured reconstruction backbone for mask-guided editing, and VoxHammer [Li et al. \(2025a\)](#) performs native 3D editing with explicit edit masks. EVA01 receives only the

Input	PointLLM-7B	ShapeLLM-7B	ShapeLLM-Omni	MiniGPT-3D	GreenPLM	EVA01 (Ours)
	This is a 3D model of a cartoon-style monkey in a rich brown color. The monkey is depicted with a long, flowing beard and is equipped with a large, drum-like hammer . The model's playful design and oversized features make it visually engaging. The monkey's posture suggests it is ready for some fun, possibly indicating it's a character from a game or animation series.	The object is a 3D model of a horse , rendered in a monochromatic white color. It showcases a detailed representation of the horse's head, including its mane and tail, and features a prominent eye.	The 3D model depicts a cartoon-like figure, resembling a man or a soldier, donning a vibrant yellow helmet. The helmet is adorned with a golden crest, suggesting a sense of regality or special distinction. In his hand, the figure holds a sword , indicating that he may be a warrior or a protector. The cartoonish style of the model, combined with the vivid colors, gives it a playful and engaging appeal, making it suitable for use in animations, video games, or other forms of digital media.	The object is a white rectangular box with black markings. It has a distinct shape and is visible from all sides. The black markings on the box are likely to be some kind of writing or symbols. The object is likely to be used as a storage container or as a decorative piece. The color of the box is	A 3D model of a knight in armor with a helmet, holding a sword and a shield , and wearing a backpack .	The object is a humanoid figure with a rugged appearance, featuring a textured surface resembling fur or scales. It has prominent horns on its head and carries a large, intricately designed sax . The body is adorned with armor-like elements, including shoulder pads and arm guards , giving it a warrior-like look. The overall color scheme includes shades of brown and gray , contributing to an ancient or mystical aesthetic.
	This is a 3D model of a robot , designed in a cartoon style. The robot is predominantly grey, suggesting a metallic composition. It has a boxy stature, reminiscent of early robot designs. The model is simplistic, lacking intricate details, but it captures the essence of a robot. Its whimsical, cartoonish style suggests it might be used in animation, video games.	The 3D model appears to be a complex assembly of interconnected mechanical parts, possibly a part of a larger system.	The 3D object model represents a toy robot designed to mimic the appearance of an AT-AT walker from Star Wars. It has a boxy, grey body with two legs , each equipped with a pair of wheels for mobility. The robot also features two guns mounted on top, which can possibly mimic the firing abilities of its Star Wars counterpart. This toy is likely used for play or as a collectible item for Star Wars fans.	This 3D model is of a white car with a red stripe. The car is likely to be a vintage model, as indicated by the red stripe. The red stripe could be a part of the car's original design or could be a result of a restoration process. The white color of the car could be a result of a thorough cleaning ...	A 3D model of a robot with two arms , two legs, and four wheels , resembling a tank.	The object is a heavily armored mech with a robust design featuring multiple barrels on its back, suggesting it's equipped for combat or heavy-duty tasks. It has articulated legs with sharp claws at the ends, indicating mobility over rough terrain. The body is covered in a camouflage pattern, giving it a tactical appearance. Various mechanical components are visible, including joints and panels that suggest functionality and flexibility.
	This is a 3D model of a ship , meticulously designed with a high level of detail. The model features a variety of parts of a ship, including masts, sails, and hulls. The realistic design suggests that it could be used for educational purposes, for example, to teach about naval architecture or maritime history. It could also serve as a decorative piece, adding a touch of sophistication and knowledge to any setting.	The 3D model appears to be a wooden ship with a prominent sail, a cabin, and a deck.	The 3D object model depicts a charming, vintage-style boat painted in a vibrant shade of blue . The boat features a mast proudly hoisting a flag at its top. A small house is attached to the side of the boat, adding character and depth to the model. This could be a playful representation of a fishing boat or a leisure boat used for recreational activities. The intricate details of the model suggest it might be used for decorative purposes or in digital animations.	This 3D model is of a white car with a red stripe. The car is likely to be a vintage model, as indicated by the red stripe. The car is also likely to be a sports car, given its sleek design. The red stripe could be a design feature or a sign of the car's high performance. The car is likely to be a collector's item due to its vintage and sports car characteristics.	3D model of a wooden pirate ship with sails and cannons, resembling the Black Pearl from Pirates of the Caribbean.	The object is a large sailing ship with multiple masts and sails. It features a complex structure with various decks, windows, and rigging. The hull has a distinct bow and stern, and there are numerous cannons along the sides. The ship's design suggests it is built for speed and maneuverability in open waters.
	This is a 3D model of a cartoon-style angel . The angel is predominantly white, with a light, airy appearance that is often associated with the supernatural beings. The model is designed with exaggerated features, including large, sweeping wings and a radiant, otherworldly light. This kind of 3D model is often used in animation, video games, or graphic design.	The 3D model appears to be a representation of a hand , possibly holding or interacting with an object, with intricate details like veins and tendons.	The 3D object model depicts a cartoonish figure of a woman adorned with a headress, featuring a large, circular design with red and white stripes. The woman's face is stylized in a playful manner, suggesting the model may be used for animation or as a toy design. The headress and the woman's attire hint at a cultural or historical context, possibly referencing Native American traditions.	This 3D model is a stylized representation of a red car with a white stripe . The car appears to be in a state of disrepair, possibly due to an accident. The white stripe could be a decorative feature.	A 3D model of a white angel statue with wings, resembling a statue of a woman holding a sword .	The object is a sculptural representation of an angelic figure with large wings, standing on a textured base that resembles clouds or waves. The figure has flowing hair and wears a garment with intricate folds. Its posture suggests a sense of grace and serenity, typical of classical angelic depictions.
	This is a 3D model of a vintage style, cartoonish house . The house is designed with exaggerated features common in cartoons, giving it a quirky, old-house-in-the-woods feel. The model is reminiscent of animated houses seen in children's shows or video games. It's a fun, creative representation of a dwelling that combines modern cartoon aesthetics.	The 3D model appears to be a wooden wagon with a canopy, adorned with various decorative elements, including a star on the front.	This is a 3D model of a vintage trolley, featuring a rustic design with a brown wooden base and a red-pink seat . The trolley has a distinctive chimney-like structure at the top, likely used for ventilation or as a decorative element. This object could be typically found in steampunk settings or as a nostalgic piece in a home decor.	This 3D model is a depiction of a woman in a white dress , standing with her arms stretched out. She is wearing a pair of black glasses, which gives the appearance of a person. The model is in detail, showcasing the woman's features such as ...	A 3D model of a small, old-fashioned wooden wagon with wheels, resembling a horse-drawn carriage.	The object is a wooden carriage with a rustic, weathered appearance. It features a cabin-like structure with a sloped roof, multiple windows, and a front-facing door. The carriage is supported by four wheels and has a hitch at the front for attachment to a horse. The wood appears aged, with visible grain and texture, and the carriage has a traditional design with a focus on functionality and simplicity.
	The 3D object model is a white, cylindrical item with a top that is open and features a distinctive circular design. It also has multiple 'legs' extending from its base.	The object is a 3D model of a futuristic spaceship , predominantly in shades of gray, with a sleek and streamlined design. The ship features a pointed nose, multiple windows, and a distinctive blue accent on its side.	The 3D object model depicts an old-fashioned, brown carriage with four wheels. The carriage is designed with intricate details, including the presence of a large, circular wheel at the front and two smaller wheels at the back. It also features a brown roof for protection against the elements, and a clear glass windowpane for visibility.	The object is a white rectangular box with a black stripe running along its length. It has a circular shape at the top and a rectangular shape at the bottom. The visible parts of the box include the black stripe and the circular shape at the top ...	A 3D model of a futuristic spaceship resembling a Star Trek Starship Enterprise with a blue and white color scheme.	The object is a large, cylindrical spacescraft with multiple sections featuring various panels, vents, and mechanical components. It has a predominantly white exterior with blue accents and several protruding elements that suggest functionality related to propulsion or habitation. The structure includes intricate detailing such as grills, hatches, and possibly weapon systems.

Figure 5. Qualitative comparison of mesh understanding results on PointLLM-200.

Table 5. Multi-turn editing evaluation. CD and PSNR evaluate unedited-region consistency; CLIP, FD_{DINOv2}, and Pref% evaluate overall editing quality.

Method	Consistency		Quality		
	CD ↓	PSNR ↑	CLIP ↑	FD ↓	Pref.% ↑
Instant3DiT* Barda et al. (2025)	0.034	21.4	24.68	265.34	0.00%
TRELLIS* Xiang et al. (2024)	0.058	18.2	26.12	241.78	2.50%
VoxHammer* Li et al. (2025a)	0.015	32.1	30.45	178.62	3.75%
Ours	0.018	29.3	70.18	89.37	93.75%

*Requires an explicit 3D edit mask at inference time; EVA01 does not use masks.

instruction and historical context at inference time.

Table 5 shows a clear separation between preservation-oriented consistency and instruction-following edit quality. VoxHammer obtains the best CD and PSNR (0.015 and 32.1), which is expected because it is given an explicit 3D edit mask and is designed to preserve unedited regions. EVA01 remains close in consistency (0.018 CD and 29.3 PSNR) despite using no mask, indicating that historical mesh tokens provide a strong implicit preservation signal. More importantly, EVA01 dominates the quality-oriented metrics: it improves CLIP from VoxHammer's 30.45 to **70.18**, reduces FD_{DINOv2} from 178.62 to **89.37**, and receives **93.75%** user preference, compared with 3.75% for VoxHammer and 2.50% for TRELLIS. This gap indicates that mask-based preservation alone is insufficient for multi-turn editing; the model must also understand the evolving object state and synthesize the requested new geometry in context.

Figures 7 and 8 illustrate the same conclusion qualitatively. In the camera, bird, and soldier examples, EVA01







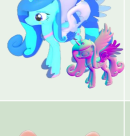
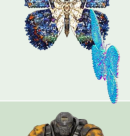
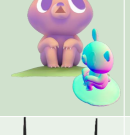
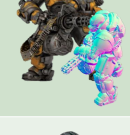
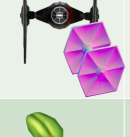

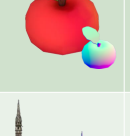

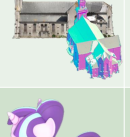
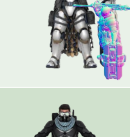
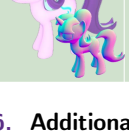

Input	GT Caption	EVA01-Alignment	EVA01-Final	Input	EVA01-Alignment	EVA01-Final
	Gray colored toy elephant with curved front teeth and long stringed tail.	a rhinoceros with horns.	The model is a detailed representation of an elephant, showcasing its large, textured body and distinctive features. It has a long trunk, large ears, and tusks. The skin appears rough and wrinkled, typical of an elephant's skin. The legs are sturdy and support the massive body, while the tail is short and slightly curved.		A small blue and yellow futuristic flying vehicle.	The model is a detailed representation of a military aircraft with a robust design. It features a blue and yellow color scheme with various markings and insignias. The aircraft has a distinctive cockpit with a transparent canopy, multiple engines mounted on the wings, and a tail section with a vertical stabilizer. The landing gear is visible, and the aircraft is equipped with various antennas and sensors. The design suggests it is built for combat or reconnaissance missions.
	White car with gray thick band running all over roof and bonnet having gray side panels and wheels	a black and white Ford Mustang GT.	The model is a detailed representation of a sports car with a sleek, aerodynamic design. It features a two-tone color scheme with a dark upper body and a lighter lower section. The car has a prominent front grille, large wheels, and a distinctive rear spoiler. The windows are tinted, and the car's body is smooth with a glossy finish. The undercarriage is visible in some views, showing a complex structure with various components.		a large, dilapidated building with a metal roof and various pipes and wires.	The model depicts a large, multi-story industrial building with a complex roof structure. The building features multiple levels with various equipment and machinery, including a prominent cylindrical tank and several interconnected sections. The roof is covered with a grid-like pattern, and there are visible pipes and conduits running across the structure. The building is surrounded by a network of roads and pathways, indicating a bustling industrial area.
	Grey coloured circular disc panel space ship with antenna.	Star Trek Starship Enterprise	The model is a futuristic spacecraft with a saucer-shaped main body and two elongated nacelles extending from the rear. The saucer has a smooth, metallic surface with a central circular feature. The nacelles are cylindrical and taper towards the ends, with a distinct separation from the main body. The overall design is sleek and streamlined, suggesting advanced technology and speed.		a rock with a hole in it.	The object is a complex, irregularly shaped structure with a rough, textured surface. It appears to be composed of intertwined, fibrous materials that give it a natural, organic appearance. The surface is predominantly a mix of earthy tones, including shades of brown, beige, and gray, with some areas showing a more uniform texture. The object has multiple protrusions and indentations, creating a dynamic and intricate form.
	A blue cartoon	a blue and white winged pony with a blue mane and tail.	The model is a blue, winged creature with a flowing mane and tail. It has large, expressive eyes and a small horn on its forehead. The wings are translucent with a gradient of blue shades, and the mane and tail have a similar color pattern. The creature's body is sleek and streamlined, with a light blue color and darker blue accents.		a blue and purple butterfly with wings.	The model is a vibrant butterfly with intricately designed wings. The wings display a rich palette of colors, including shades of blue, red, and orange, with detailed patterns resembling feathers. The body is adorned with a series of decorative elements, including a prominent circular ornament and a series of small, colorful dots. The wings have a layered appearance, with each layer contributing to the overall complexity and beauty of the design.
	A cartoon small dog pink	a purple and pink teddy bear on a green and purple base.	The model is a stylized, cartoonish mouse sitting on a green circular base. It has a light purple body with a darker purple belly and a small tail. The mouse's face features large, round eyes and a small nose. Its ears are small and rounded, and it has a playful expression. The base is simple and smooth, providing a stable platform for the mouse.		a yellow and black armored robot with a gun.	The model is a heavily armored figure with a futuristic design, featuring a combination of yellow and black colors. It has a robust, bulky appearance with intricate details such as a large backpack, a helmet with a visor, and various mechanical components. The armor is segmented, with visible joints and panels, suggesting flexibility and protection. The figure is equipped with a large weapon, possibly a gun, and has additional gear attached to its back, indicating a combat-ready or tactical role.
	Grey fin threaded hexagonal umbrella panel sided camera.	Star Wars TIE Fighter	The model is a spacecraft with a central spherical cockpit featuring a red circular window. It has two large, hexagonal solar panels on either side, each with a grid-like pattern. The panels are connected to the cockpit by a central support structure. The overall design is sleek and symmetrical, with a metallic texture and a dark color scheme.		a black and white motorcycle with a large wheel and a black seat.	The model is a detailed representation of a motorcycle with a sleek, aerodynamic design. It features a prominent headlight, a streamlined body, and a large, textured engine. The wheels are robust, and the overall color scheme is a mix of dark and metallic tones, giving it a modern and sophisticated appearance.
	A cartoon green and red like a fruit	Low-poly red apple with a green leaf.	The model is a stylized representation of a red apple with a smooth, rounded surface. It features a green leaf attached to a short brown stem. The apple has a consistent red color with slight variations in shading, giving it a realistic appearance. The leaf is broad and flat, extending from the top of the apple. The stem is short and slightly curved, emerging from the top center of the apple.		a yellow and black armored truck.	The model is a rugged, yellow off-road vehicle with a robust design. It features large, thick tires suitable for rough terrain, a sturdy frame, and a cabin with a protective windshield. The vehicle is equipped with various functional components, including a spare tire mounted on the back, a roof rack, and a visible exhaust system. The vehicle's exterior is marked with a distinctive logo, and it has a utilitarian appearance, suggesting it is designed for heavy-duty use.
	A cartoon green look like a very big and old church	a church with a steeple.	The model is a detailed representation of a historic church with a tall, pointed spire. The structure features a large, arched entrance with a stone facade and multiple windows. The roof is steeply pitched, and the building is surrounded by a grassy area. The spire is ornate, with a cross at the top, and the walls are constructed from stone, giving the building a robust and ancient appearance.		a white and black armored soldier with a sword, holding a gun and a shield.	The model is a detailed representation of a futuristic soldier in full armor. The armor is composed of various segments, including a helmet, chest plate, shoulder guards, and leg armor. The soldier is equipped with a long, slender weapon resembling a rifle. The armor features a mix of metallic and matte finishes, with intricate detailing on the chest and shoulder areas. The helmet has a visor, and the overall design suggests a focus on protection and mobility.
	A cartoon pony	a purple and white pony with a long tail.	The model is a stylized unicorn with a light purple body and a darker purple mane and tail. It has a white face with large, expressive eyes and a small horn on its forehead. The mane and tail are voluminous and flow down the back. The unicorn's legs are sturdy, and it has a small, rounded tail. The overall appearance is smooth and cartoonish, with a playful and friendly demeanor.		a man in a suit with a gun	The model is a humanoid figure with a skeletal appearance, featuring a skull-like head and a long, skeletal arm. It is dressed in tattered, dark clothing with a belt and a large, mechanical-looking weapon strapped to its back. The figure has a menacing expression and appears to be in a dynamic, action-oriented pose, suggesting movement or combat readiness.

Figure 6. Additional Mesh Understanding Examples. Left: randomly sampled PointLLM-200 examples with the rendered input, official reference caption, EVA01-Align caption, and EVA01-Final caption. Right: captions for generated 3D models without ground-truth annotations. EVA01-Final provides richer part-, material-, color-, and structure-level descriptions than the alignment-stage model.

accumulates edits across three rounds while preserving the source identity: added parts remain attached in later turns, removed components do not reappear, and pose or articulation changes are applied without destroying previously edited structure. The additional long-horizon trajectories further stress harder edits, including replacing a knight's shield before adding a cloak, attaching robotic arms and a roof cannon to a rover, transforming a soldier into a rider on a mechanical dinosaur, and repeatedly modifying a garage with a vault hatch, propeller, satellite dish, thrusters, and reinforced exterior panels. These examples require constructive geometry, part replacement, articulation, accessory insertion, and material preservation over two to five turns. EVA01's ability to keep RGB and normal renderings aligned across these trajectories suggests that it maintains a coherent latent state rather than merely producing visually plausible single-step outputs.

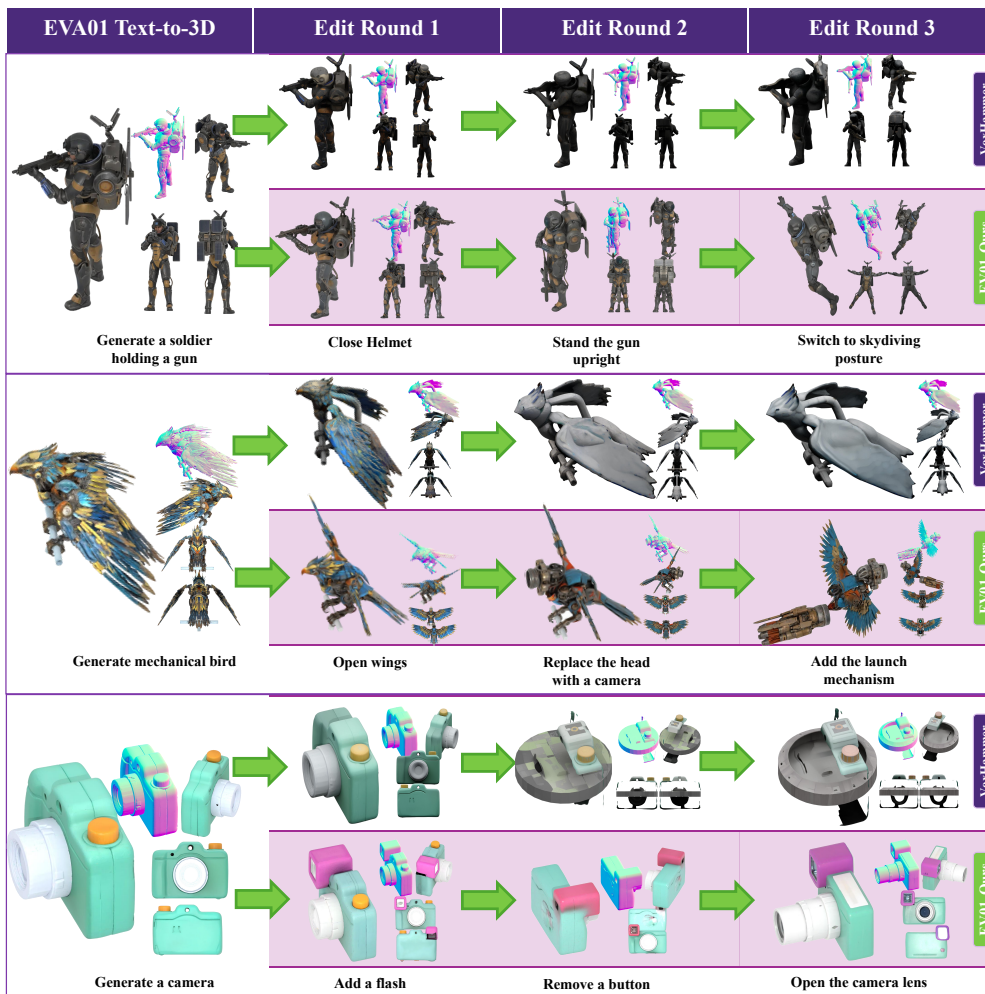


Figure 7. Versatile Instruction-Driven 3D Editing Gallery. Starting from EVA01 text-to-3D generations, we compare three rounds of sequential edits against VoxHammer. Across a camera, a mechanical bird, and a soldier, EVA01 accumulates instructions—adding or removing parts, changing object state, replacing components, and modifying pose—while preserving object identity and geometry history. Additional long-horizon trajectories are shown in Figure 8.

4.5 Subjective Evaluation

Study Protocol. We conduct user studies across all three evaluation axes: single-turn text-to-3D generation, single-turn image-to-3D generation, and multi-turn editing. A total of 23 participants are recruited; each completes at least 25 paired or multi-way comparisons per category. For each comparison, participants are shown the input condition (text prompt, reference image, or editing instruction with interaction history) alongside anonymized outputs from competing methods, presented in randomized order. Participants are instructed to select the output that best satisfies the given condition, judging geometric fidelity, semantic alignment, and—for editing tasks—identity preservation across turns. All selections are aggregated per category, and preference scores are reported as the percentage of total selections each method receives within each category.

Results. Table 3 reports preference results for single-turn generation. In text-to-3D, EVA01 receives **70.4%** of all selections, a margin of nearly $5\times$ over TRELIS (14.8%) and over $8\times$ over the text-specialized 3DGen-R1 (7.9%). In image-to-3D, TRELIS.2 leads with 41.7%; EVA01 places second at **22.2%**, ahead of Hunyuan3D-2.1 (10.2%), Step1X-3D (7.9%), and other supported baselines. Notably, EVA01 is the only method to rank among the top two in both modalities within a single unified model.

Table 5 reports preference for multi-turn editing. EVA01 receives **93.75%** of selections, compared to 3.75% for VoxHammer and 2.50% for TRELIS, despite both baselines having access to explicit 3D edit masks while EVA01 operates mask-free. This large preference gap reflects the advantage of stateful editing: EVA01 conditions each edit on the full interaction history, preserving geometric identity across turns, whereas mask-based baselines treat each edit as an independent reconstruction. Participants consistently preferred EVA01’s outputs for maintaining object identity while applying the requested structural changes across sequential editing instructions.




Figure 8. Additional Generation and Editing Results. Each cell shows a trajectory state with a main render, auxiliary RGB/normal views, and a concise generation or editing prompt.

4.6 Representation Analysis and Visualization

Following recent protocols for probing 3D awareness in foundation models [El Banani et al. \(2024\)](#); [Huang et al. \(2025\)](#), Table 6 and Figure 9 evaluate which feature paths can provide visual, semantic, or generative conditioning for a 3D-native MLLM. NAVI R@5cm and the NYUv2 depth probes measure dense spatial and geometric information: NAVI tests whether cross-view correspondences fall within a 5cm 3D threshold, while NYUv2 δ_1 and RMSE measure how linearly scene layout, depth ordering, and geometric boundaries can be read from frozen features. Objaverse retrieval R@5 follows a different protocol: image and caption features are mean-pooled, normalized, and matched by top-5 paired-caption retrieval. For contrastive models such as CLIP and SigLIP2, this directly probes the native image-text embedding space; for Qwen3-VL, Bagel, and Qwen-Image-Edit, it is a narrower test of whether hidden states can be pooled into retrieval embeddings, rather than a general measure of semantic understanding.

Table 6. Feature probing across visual and image-generation encoders. Higher values are better for NAVI R@5cm, NYUv2 δ_1 , and Objaverse retrieval R@5; lower values are better for NYUv2 RMSE. Bold and underline indicate the best and second-best entries within each category block. Lavender marks semantic/understanding-token paths; mint marks dense visual or generation-latent paths. “-” denotes metrics not applicable to raw visual/latent-only features in the retrieval protocol.

Model	Feature path	NAVI R@5cm \uparrow	NYUv2 $\delta_1\uparrow$	NYUv2 RMSE \downarrow	Retrieval R@5 \uparrow
Visual baselines					
CLIP ViT-L/14-336	final	47.96	0.3859	1.1866	43.50
SigLIP2 SO400M	final	78.30	0.6659	0.6567	<u>23.00</u>
DINOv2 ViT-L/14	final	<u>91.85</u>	<u>0.8772</u>	<u>0.3341</u>	-
DINOv3 ViT-L/16	final	92.74	0.8935	0.3126	-
Qwen3-VL DeepStack / LLM					
Qwen3-VL-2B	VT ds5	66.04	0.4797	0.9802	-
Qwen3-VL-2B	VT ds17	<u>78.37</u>	<u>0.7510</u>	0.5612	-
Qwen3-VL-2B	VT final	69.54	0.6897	0.6308	-
Qwen3-VL-2B	LLM l07	75.23	0.7439	<u>0.5460</u>	0.50
Qwen3-VL-2B	LLM l14	73.88	0.7826	0.5246	0.50
Qwen3-VL-2B	LLM l28	71.77	0.7396	0.5868	15.80
Qwen3-VL-4B	VT ds17	77.79	0.6842	0.6371	-
Qwen3-VL-4B	LLM l14	75.40	0.7458	0.5569	0.50
Qwen3-VL-8B	VT ds16	80.25	0.7155	0.5975	-
Qwen3-VL-8B	LLM l07	76.68	0.7474	0.5542	0.50
Image-generation MLLMs					
Bagel-7B-MoT	ViT final	76.45	0.6320	0.7009	-
Bagel-7B-MoT	VAE clean	53.56	0.3929	1.1725	-
Bagel-7B-MoT	VAE -> LLM l21	61.11	<u>0.5923</u>	<u>0.7249</u>	0.60
Qwen-Image-Edit	Qwen2.5-VL final	62.21	0.5901	0.7620	17.50
Qwen-Image-Edit	MMDiT 20/30/40	<u>68.51</u>	0.5509	0.8318	0.60
Qwen-Image-Edit	Ref VAE latent	55.04	0.3896	1.1877	-

The visual baselines separate global semantic alignment from dense geometric structure. CLIP obtains the highest retrieval score (43.50), but weak NAVI and NYUv2 results (47.96, 0.3859, and 1.1866). Figure 9 explains this gap: CLIP’s PCA maps capture coarse object-level regions and its self-similarity maps form high-contrast semantic blobs, whereas its norm maps are sparse and edge-biased, and its cross-view matches are less consistently anchored to local surface geometry. Thus, global image-text contrastive alignment does not by itself provide a stable correspondence field. In contrast, DINO-style dense features define the geometric upper envelope. DINOv3 ViT-L/16 achieves the strongest NAVI, δ_1 , and RMSE scores (92.74, 0.8935, and 0.3126); its PCA and norm maps are spatially smooth, object- and part-aligned, and preserve coherent room layout and object boundaries across Objaverse, NYUv2, and NAVI examples, with more geometrically plausible cross-view matches. DINOv2 shows the same dense-visual behavior with slightly weaker boundary and layout separation. SigLIP2 lies between the two regimes: although optimized as an image-text model, its maps collapse less than CLIP’s and retain more local structure, consistent with its stronger NAVI and NYUv2 scores despite lower retrieval.

Qwen3-VL shows that 3D-relevant information is distributed across paths and layers rather than concentrated in the final hidden state. Along the visual-token path, middle-to-late DeepStack features preserve stronger local geometry: Qwen3-VL-8B VT ds16 gives the best NAVI score in the Qwen3-VL block (80.25), and Qwen3-VL-2B VT ds17 is close (78.37), both clearly above early VT ds5 and the visual-encoder final state. The visualization shows the same pattern: DeepStack PCA maps retain object and part regions, and cross-view matches remain tied to corresponding local surfaces, whereas early VT features are smoother and less discriminative, and the final visual state is more mixed. LLM image-token layers shift the representation toward scene-level reasoning. Qwen3-VL-2B LLM l14 gives the strongest NYUv2 result in the Qwen3-VL group (0.7826 δ_1 , 0.5246 RMSE), and its NYUv2 PCA and self-similarity panels emphasize broader room-layout regions rather than fine object parts. Later LLM layers move further toward global semantic compatibility: Qwen3-VL-2B LLM l28 has the highest retrieval score in this group (15.80), while its dense probing scores weaken. These trends favor layer- and path-aware routing over simply using the last hidden state or scaling the model.

Image-generation MLLMs exhibit a complementary limitation. In Bagel, the ViT semantic path is the strongest



Figure 9. Representation visualization across visual and image-generation encoders. We visualize the same feature paths probed in Table 6 using input views, PCA projections, normalized activation maps, self-similarity, and cross-view correspondence overlays. Lavender rows denote semantic or understanding-token paths, while mint rows denote dense visual or generation-side latent paths. The visualization reveals that global semantic alignment, dense spatial correspondence, and generation-side appearance latents form distinct representation regimes rather than a single universally optimal feature space.

feature for dense probing, while the clean VAE latent is much weaker. The figure makes this distinction explicit: clean VAE latents preserve appearance, reconstruction layout, and high-frequency edges in the PCA, norm, and self-similarity maps, especially on synthetic Objaverse objects, but this reconstruction-oriented signal does not translate into robust NAVI or NYUv2 geometry. Passing VAE latents through the LLM partially restores higher-level structure and improves probing scores, yet remains below the best visual and Qwen3-VL paths. Qwen-Image-Edit shows the same division of labor. Its Qwen2.5-VL final state is better suited to semantic compatibility and retrieval, whereas MMDiT blocks expose denser local cues: their PCA and self-similarity maps preserve object silhouettes, edges, and scene layout more clearly than the purely semantic path. The reference VAE latent again behaves as an appearance/reconstruction condition, with crisp contours but weak transferable geometric abstraction. Generation-side latents are therefore useful for synthesis and appearance anchoring, but should not be treated as complete 3D geometry representations.

These results motivate EVA01’s separation of semantic reasoning from geometric generation. A 3D-native system cannot reduce visual conditioning to a single frozen encoder, final MLLM hidden state, or VAE latent. Table 6 shows that semantic retrieval, cross-view correspondence, and depth probing peak in different feature families; Figure 9 shows the same split through coarse semantic blobs, dense object-aligned PCA fields, layout-sensitive self-similarity maps, and cross-view match stability. EVA01 therefore keeps the Understanding Expert as a stable semantic anchor while allowing the Generation Expert to access both semantic tokens and dense visual-geometric cues through shared global attention. The CLIP–DINO split, Qwen3-VL’s DeepStack/LLM specialization, and the Bagel/Qwen-Image-Edit separation among VAE, transformer, and semantic paths all support the same design principle: semantic reasoning and geometric generation should be decoupled, yet remain communicative within a shared sequence space.

4.7 Ablation Studies and Critical Insights

Figure 10 summarizes the ablations that shaped EVA01’s final curriculum and architecture. The central question is not only whether a mesh representation can be decoded into geometry, but whether it can be aligned with the pre-trained MLLM semantic space, optimized inside a unified sequence, and reused for context-aware generation without destroying the language prior.

The mesh-understanding curves reveal that alignment must precede instruction tuning. Directly finetuning on mesh captions lowers CE loss rapidly in the first few thousand steps, but the curve soon saturates and reaches a weaker normalized captioning score. In contrast, the 10K alignment warm-up descends more slowly at the beginning, yet it creates a cleaner semantic bridge between Point-BERT mesh features and the Qwen3-VL token space. Once instruction tuning starts, this aligned model overtakes the direct-tuning baseline and converges to both lower CE loss and higher captioning quality. The Sparse Shape-to-Text variant exposes a complementary failure mode: directly using generation-side sparse VAE latents as the mesh-understanding input, even when the generation pathway is treated as an encoder and fully tuned, does not yield a stable captioning model. Its loss plateaus around a high CE regime and the generated descriptions remain semantically unreliable. This confirms that sparse generation latents are effective reconstruction variables, but they are not by themselves language-readable semantic tokens.

The generation curves show why the curriculum begins from image-conditioned generation before text-to-3D alignment. Training only from text reduces MSE, but it converges more slowly and reaches a lower score. Image warm-up provides a stronger bridge because dense visual features already live closer to the MLLM’s semantic manifold and carry local geometric evidence through the DeepStack pathway. Adding mesh-understanding samples further improves both convergence and the final score, indicating that the captioning objective is not merely an auxiliary regularizer: it sharpens the spatial grounding of E_{und} , and shared global attention allows E_{gen} to query these better-aligned semantic features during generation. This bidirectional coupling is also more effective than directly feeding multi-layer hidden features through concatenation-based cross-attention, which supplies additional conditioning but lacks token-level mutual visibility inside the unified sequence. As a result, it improves over text-only training but remains below the MoT formulation in both efficiency and upper bound, and it does not naturally support context-aware editing over interleaved mesh histories.

The same figure also explains why EVA01 adopts a structured sparse grid rather than a permutation-invariant VecSet representation [Zhang et al. \(2023\)](#), including the variant used in Hunyuan3D-2.1 [Team \(2025a\)](#). Under our unified sequence setting, VecSet fails to produce usable geometry: the loss quickly reaches a plateau and the normalized score remains far below grid-based sparse latents. Although VecSet is compact, its latent tokens do not carry an intrinsic position anchor after being flattened into the MLLM sequence. Attention therefore observes a weak global token identity, but cannot reliably distinguish where each latent token is located in 3D space. In contrast, our sparse voxel tokens bind each feature to an explicit coordinate p_i , making 3D Interleaved MRoPE meaningful and allowing attention to model local and global spatial relations. These results support three practical design rules for 3D-native MLLMs: modality dropout is needed to prevent E_{gen} from ignoring weak text conditions during Stage 3; Mesh-to-Text supervision provides a necessary bidirectional alignment signal rather than a secondary task; and high-quality finetuning in Stage 5 sets the final fidelity ceiling after the representation and alignment problems have been stabilized.

5 Limitations, Discussion & Future Work

EVA01 is bounded by the training budget of this study: both experts operate at the 2B scale with 512^3 sparse-voxel resolution. This setting suffices to verify the central design hypothesis—that semantic reasoning and geometric generation should be decoupled yet connected through shared global attention—but does not exhaust the scaling behavior of 3D-native MLLMs. Scaling each expert toward 4B–8B and increasing resolution to 1024^3 is a natural next step, where larger context and finer spatial discretization are expected to improve both long-context reasoning and high-frequency geometric recovery.

Current automatic 3D generation metrics remain partial proxies for perceptual and geometric quality. Metrics built on aligned 3D representation spaces, such as ULIP, ULIP-2, and Uni3D [Xue et al. \(2023, 2024\)](#); [Zhou et al. \(2024\)](#), measure coarse semantic agreement between language, images, and 3D shapes, but operate on sampled point clouds or global shape embeddings. They under-resolve surface-level properties critical for high-fidelity mesh generation: local topology, thin structures, sharp creases, watertightness, material-boundary consistency, and whether an edit preserves the identity of unmodified regions. This creates a mismatch between improving generators and comparatively coarse evaluators. Developing benchmarks for fine-grained mesh quality, semantic

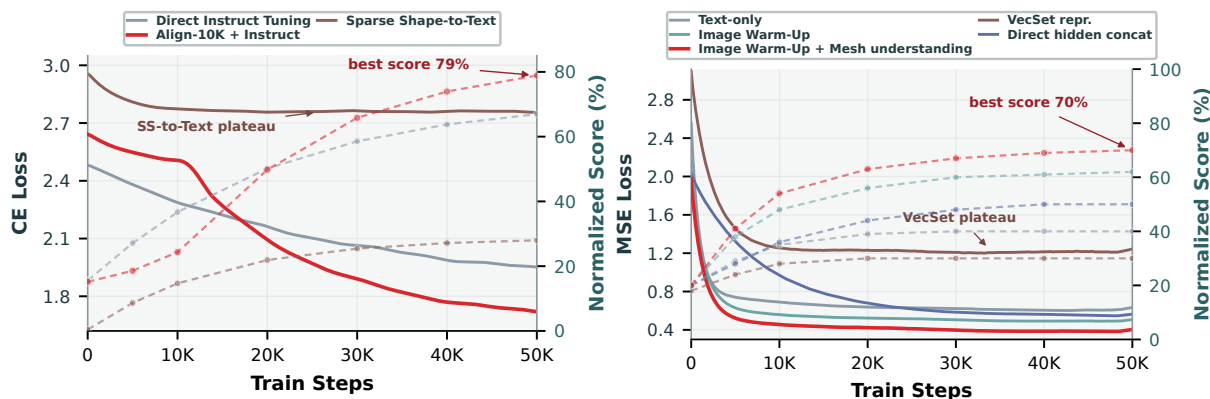


Figure 10. Training Dynamics and Loss Curves. Left: mesh-understanding ablations comparing direct instruction tuning, a 10K alignment warm-up followed by instruction tuning, and Sparse Shape-to-Text, which uses generation-side sparse VAE latents for captioning. Solid curves report CE loss, and dashed marker curves report normalized captioning score. Right: Sparse Shape generation ablations comparing text-only training, image warm-up, image warm-up with mesh understanding, VecSet representation, and multi-layer hidden-feature concatenation for cross-attention. Solid curves report MSE loss, and dashed marker curves report normalized generation score relative to the best text-to-3D checkpoint.

Text-to-3D

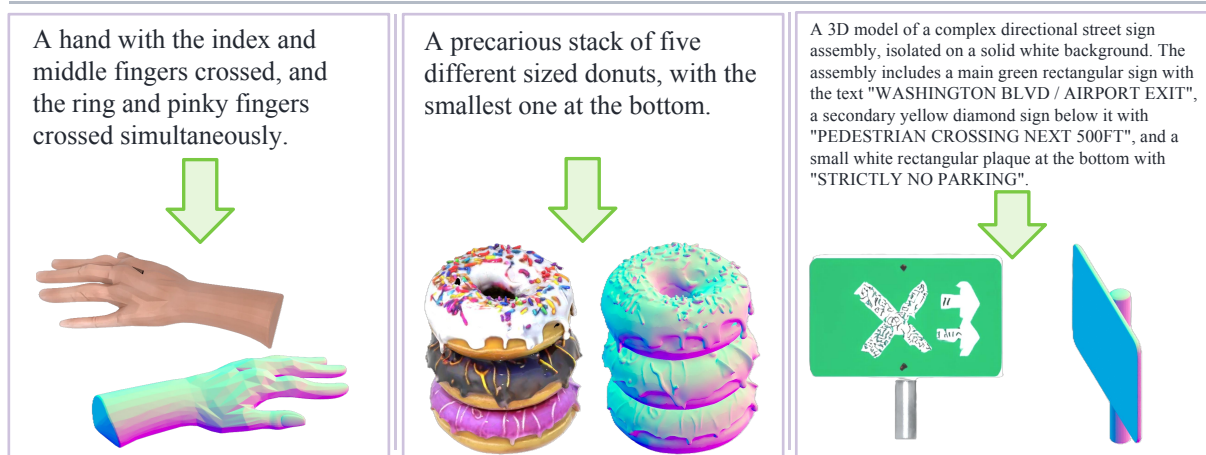


Image-to-3D




Figure 11. Failure Cases of EVA01. In text-to-3D generation, EVA01 still has limited generalization to out-of-distribution compositions and remains imperfect at spatial reasoning, exact counting, and producing legible text or symbol layouts on 3D surfaces. In image-to-3D cases, failure is often caused by insufficient dense pixel evidence in the input view: when thin structures, occluded parts, or small distant components are under-resolved, the model may recover the dominant object while losing local details or producing incomplete geometry.

faithfulness, and context-aware editing consistency remains an important direction.

A third limitation is the representational asymmetry between mesh understanding and mesh generation. Even architectures that decouple visual encoders—Janus and Bagel Wu et al. (2024); Deng et al. (2025)—consume RGB images on both sides. In EVA01, the understanding branch relies on point-cloud features, while the generation branch operates on structured sparse-voxel latents. This split is pragmatic: point-cloud encoders provide a stable entry point for mesh captioning, whereas sparse voxels are better matched to high-resolution geometry synthesis and 3D positional attention. A more native interface would replace this split with a shared 3D latent substrate. Inspired by unified visual encoder families such as OpenVision and OpenVision 3 Li et al. (2025c); Zhang et al. (2026), future work should explore sparse-voxel-native encoder families in which understanding and generation share a common 3D representation served by specialized encoders over the same latent backbone.

The image-conditioned setting exposes another boundary, with representative failure modes shown in Figure 11. DeepStack injects lower-level SigLIP2 features into the MLLM backbone, but our representation analysis shows these features are weaker dense geometric carriers than DINO-style self-supervised visual features. Consequently, EVA01 is competitive in image-to-3D while trailing the strongest image-specialized reconstruction pipeline. This does not reflect a failure of the MoT formulation; rather, a 3D-native MLLM requires sufficiently local, patch-aligned visual evidence for the Generation Expert to reconstruct fine geometry from a single image. Replacing the current SigLIP2 visual path with encoders that jointly preserve text alignment and dense patch structure—such as TIPSv2 Cao et al. (2026)—or adopting a multi-tower visual aggregation design in the spirit of Cambrian-1’s Spatial Vision Aggregator Tong et al. (2024) could strengthen the pixel-dense conditioning available to E_{gen} without sacrificing semantic grounding.

Taken together, these limitations delineate the scope of this work. EVA01 does not claim that a 2B, 512^3 instance saturates all 3D generation regimes, nor that existing metrics fully capture mesh quality. Its contribution is a scalable architectural template: a semantic expert that preserves MLLM priors, a generation expert that learns sparse geometric flow matching, and a shared attention interface through which semantic and geometric tokens remain communicative. The gaps identified here—larger model scale, higher spatial resolution, finer evaluators, unified 3D-native encoders, and stronger dense visual conditioning—define a concrete path toward the next generation of 3D-native multimodal foundation models.

6 Conclusion

We presented **EVA01**, a unified framework that integrates 3D mesh understanding, generation, and multi-turn editing within a single Mixture-of-Transformers architecture. By decoupling semantic understanding from geometric generation via a dual-expert design with shared global attention, EVA01 transfers multimodal priors from a pre-trained MLLM backbone to the 3D domain, bridging the semantic-geometric alignment gap under limited 3D supervision.

EVA01 achieves state-of-the-art native text-to-3D generation fidelity and enables context-aware, identity-preserving multi-turn 3D editing—a capability inaccessible to stateless reconstruction pipelines. Our experiments yield three practical design principles for 3D-native MLLMs: grid-based sparse latents are necessary for geometric validity under unified sequence modeling; modality dropout and multi-stage curriculum training are essential for bridging textual and geometric manifolds; and high-fidelity finetuning sets the final fidelity ceiling after alignment is stabilized.

These findings establish a scalable architectural template for 3D-native multimodal models. The limitations discussed in Section 5—model scale, spatial resolution, evaluation fidelity, unified 3D-native encoders, and dense visual conditioning—define natural directions for future work.

Authors¹

Team Leaders. Zhengdong Guo; Shimu Wang. **Algorithm Leader.** Zongyuan Yang.

Core Contributors. Wanli Ma; Zongyuan Yang; Mingjing Yi.

Contributors. Chenzhuo Fan; Bocheng Li; Baolin Liu; Yuke Lou; Yingde Song; Qianchi Yang.

¹Alphabetical by last name within each role.

References

- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, Mei Li, Kaixin Li, Zicheng Lin, Junyang Lin, Xuejing Liu, Jiawei Liu, Chenglong Liu, Yang Liu, Dayiheng Liu, Shixuan Liu, Dunjie Lu, Ruilin Luo, Chenxu Lv, Rui Men, Lingchen Meng, Xuancheng Ren, Xingzhang Ren, Sibao Song, Yuchong Sun, Jun Tang, Jianhong Tu, Jianqiang Wan, Peng Wang, Pengfei Wang, Qiuyue Wang, Yuxuan Wang, Tianbao Xie, Yiheng Xu, Haiyang Xu, Jin Xu, Zhibo Yang, Mingkun Yang, Jianxin Yang, An Yang, Bowen Yu, Fei Zhang, Hang Zhang, Xi Zhang, Bo Zheng, Humen Zhong, Jingren Zhou, Fan Zhou, Jing Zhou, Yuanzhi Zhu, and Ke Zhu. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*, 2025.
- Satanjeev Banerjee and Alon Lavie. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan, 2005. Association for Computational Linguistics.
- Amir Barda, Matheus Gadelha, Vladimir G Kim, Noam Aigerman, Amit H Bermano, and Thibault Groueix. Instant3DiT: Multiview inpainting for fast editing of 3D objects. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 16273–16282, 2025.
- Bingyi Cao, Koert Chen, Kevis-Kokitsi Maninis, Kaifeng Chen, Arjun Karpur, Ye Xia, Sahil Dua, Tanmaya Dabral, Guangxing Han, Bohyung Han, et al. Tipsv2: Advancing vision-language pretraining with enhanced patch-text alignment. *arXiv preprint arXiv:2604.12012*, 2026.
- Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. Shapenet: An information-rich 3d model repository. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1918–1927, 2015. URL <https://arxiv.org/abs/1512.03012>.
- Wenyue Chen, Wenjue Chen, Peng Li, Qinghe Wang, Xu Jia, Heliang Zheng, Rongfei Jia, Yuan Liu, and Ronggang Wang. Know3d: Prompting 3d generation with knowledge from vision-language models. *arXiv preprint arXiv:2603.22782*, 2026.
- Xiaokang Chen, Chengyue Wu, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, Chong Ruan, and Ping Luo. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint*, 2025a.
- Yongwei Chen, Yushi Lan, Shangchen Zhou, Tengfei Wang, and Xingang Pan. Sar3d: Autoregressive 3d object generation and understanding via multi-scale 3d vqvae. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 28371–28382, 2025b.
- Zhaoxi Chen, Jiayang Tang, Yuhao Dong, Ziang Cao, Fangzhou Hong, Yushi Lan, Tengfei Wang, Haozhe Xie, Tong Wu, Shunsuke Saito, et al. 3DTopia-XL: Scaling high-quality 3D asset generation via primitive diffusion. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 26576–26586, 2025c.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53, 2024.
- Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers, 2023. URL <https://arxiv.org/abs/2309.16588>.
- Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects, 2022. URL <https://arxiv.org/abs/2212.08051>.
- Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, Eli VanderBilt, Aniruddha Kembhavi, Carl Vondrick, Georgia Gkioxari, Kiana Ehsani, Ludwig Schmidt, and Ali Farhadi. Objaverse-xl: A universe of 10m+ 3d objects, 2023. URL <https://arxiv.org/abs/2307.05663>.

- Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, Guang Shi, and Haoqi Fan. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*, 2025.
- Runpei Dong, Chunrui Han, Yuang Peng, Zekun Qi, Zheng Ge, Jinrong Yang, Liang Zhao, Jianjian Sun, Hongyu Zhou, Haoran Wei, et al. Dreamllm: Synergistic multimodal comprehension and creation. In *Proceedings of ICLR*, 2024.
- Mohamed El Banani, Amit Raj, Kevis-Kokitsi Maninis, Abhishek Kar, Yuanzhen Li, Michael Rubinstein, Deqing Sun, Leonidas Guibas, Justin Johnson, and Varun Jampani. Probing the 3D awareness of visual foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21795–21806, 2024.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. SimCSE: Simple contrastive learning of sentence embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic, 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.552.
- Xiaoguang Han, Yushuang Wu, Luyue Shi, Haolin Liu, Hongjie Liao, Lingteng Qiu, Weihao Yuan, Xiaodong Gu, Zilong Dong, and Shuguang Cui. Mvimngnet 2.0: A larger-scale dataset of multi-view images. *arXiv preprint*, 2024.
- Xianglong He, Junyi Chen, Sida Peng, Di Huang, Yangguang Li, Xiaoshui Huang, Chun Yuan, Wanli Ouyang, and Tong He. GVGEN: Text-to-3D generation with volumetric representation. In *European Conference on Computer Vision*, pages 463–479. Springer, 2024.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. CLIPScore: a reference-free evaluation metric for image captioning. In *EMNLP*, 2021.
- Junming Huang and Weiwei Xu. Cg-mllm: Captioning and generating 3d content via multi-modal large language models. *arXiv preprint arXiv:2601.21798*, 2026.
- Peng Huang, Yifeng Chen, Zeyu Zhang, and Hao Tang. Unimesh: Unifying 3d mesh understanding and generation. *arXiv preprint arXiv:2604.17472*, 2026.
- Zixuan Huang, Xiang Li, Zhaoyang Lv, and James M Rehg. How much 3D do video foundation models encode? *arXiv preprint arXiv:2512.19949*, 2025.
- Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. In *International conference on machine learning*, pages 4651–4664. PMLR, 2021.
- Tanghai Jia, Dongyu Yan, Dehao Hao, Yang Li, Kaiyi Zhang, Xianyi He, Lanjiong Li, Jinnan Chen, Lutao Jiang, Qishen Yin, Long Quan, Ying-Cong Chen, and Li Yuan. Ultrashape 1.0: High-fidelity 3d shape generation via scalable geometric refinement. *arXiv preprint arXiv:2512.21185*, 2025.
- Heewoo Jun and Alex Nichol. Shap-E: Generating conditional 3D implicit functions. *arXiv preprint arXiv:2305.02463*, 2023.
- Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the Fourth Eurographics Symposium on Geometry Processing (SGP)*, pages 61–70, 2006.
- Zeqiang Lai, Yunfei Zhao, Zibo Zhao, Haolin Liu, Qingxiang Lin, Jingwei Huang, Chunchao Guo, and Xiangyu Yue. Lattice: Democratize high-fidelity 3d generation at scale, 2025. URL <https://arxiv.org/abs/2512.03052>.
- Lin Li, Zehuan Huang, Haoran Feng, Gengxiong Zhuang, Rui Chen, Chunchao Guo, and Lu Sheng. Voxhammer: Training-free precise and coherent 3D editing in native 3D space. *arXiv preprint arXiv:2508.19247*, 2025a.
- Weiyu Li, Xuanyang Zhang, Zheng Sun, Di Qi, Hao Li, Wei Cheng, Weiwei Cai, Shihao Wu, Jiarui Liu, Zihao Wang, et al. Step1X-3D: Towards high-fidelity and controllable generation of textured 3D assets. *arXiv preprint arXiv:2505.07747*, 2025b.
- Xianhang Li, Yanqing Liu, Haoqin Tu, and Cihang Xie. Openvision: A fully-open, cost-effective family of advanced vision encoders for multimodal learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3977–3987, 2025c.

- Weixin Liang, LILI YU, Liang Luo, Srini Iyer, Ning Dong, Chunting Zhou, Gargi Ghosh, Mike Lewis, Wen tau Yih, Luke Zettlemoyer, and Xi Victoria Lin. Mixture-of-transformers: A sparse and scalable architecture for multi-modal foundation models. *Transactions on Machine Learning Research*, 2025. ISSN 2835-8856. URL <https://openreview.net/forum?id=Nu6N69i8SB>.
- Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain, 2004. Association for Computational Linguistics.
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *11th International Conference on Learning Representations, ICLR 2023*, 2023.
- William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM SIGGRAPH Computer Graphics*, 21(4):163–169, 1987. doi: 10.1145/37402.37422.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.
- Yiyang Ma, Xingchao Liu, Xiaokang Chen, Wen Liu, Chengyue Wu, Zhiyu Wu, Zizheng Pan, Zhenda Xie, Haowei Zhang, Xingkai Yu, et al. Janusflow: Harmonizing autoregression and rectified flow for unified multimodal understanding and generation. *arXiv preprint*, 2024.
- Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Hervé Jégou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision, 2023. URL <https://arxiv.org/abs/2304.07193>.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. BLEU: A method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, 2002. Association for Computational Linguistics.
- William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of ICCV*, 2023.
- Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *ICLR*, 2023.
- Zekun Qi, Runpei Dong, Shaochen Zhang, Haoran Geng, Chunrui Han, Zheng Ge, Li Yi, and Kaisheng Ma. Shapellm: Universal 3d object understanding for embodied interaction. In *European Conference on Computer Vision*, pages 214–238. Springer, 2024.
- Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, pages 3982–3992, 2019.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of CVPR*, pages 10684–10695, 2022.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, Francisco Massa, Daniel Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana, Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé Jégou, Patrick Labatut, and Piotr Bojanowski. DINOv3, 2025. URL <https://arxiv.org/abs/2508.10104>.
- Stefan Stojanov, Anh Thai, and James M. Rehg. Using shape to categorize: Low-shot learning with an explicit shape bias. 2021.

- Yiwen Tang, Zoey Guo, Kaixin Zhu, Ray Zhang, Qizhi Chen, Dongzhi Jiang, Junli Liu, Bohan Zeng, Haoming Song, Delin Qu, et al. Are we ready for RL in text-to-3D generation? a progressive investigation. *arXiv preprint arXiv:2512.10949*, 2025a.
- Yuan Tang, Xu Han, Xianzhi Li, Qiao Yu, Yixue Hao, Long Hu, and Min Chen. Minigt-3d: Efficiently aligning 3d point clouds with large language models using 2d priors. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 6617–6626. Association for Computing Machinery, 2024.
- Yuan Tang, Xu Han, Xianzhi Li, Qiao Yu, Jinfeng Xu, Yixue Hao, Long Hu, and Min Chen. More text, less point: Towards 3d data-efficient point-language understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 7284–7292, 2025b.
- Tencent Hunyuan3D Team. Hunyuan3d 1.0: A unified framework for text-to-3d and image-to-3d generation, 2024.
- Tencent Hunyuan3D Team. Hunyuan3d 2.1: From images to high-fidelity 3d assets with production-ready pbr material, 2025a.
- Tencent Hunyuan3D Team. Hunyuan3d 2.0: Scaling diffusion models for high resolution textured 3d assets generation, 2025b.
- Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai C Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, et al. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *Advances in Neural Information Processing Systems*, 37:87310–87356, 2024.
- Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, et al. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv preprint*, 2025.
- Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/1a87980b9853e84dfb295855b425c262-Abstract-Conference.html.
- Zhengyi Wang, Jonathan Lorraine, Yikai Wang, Hang Su, Jun Zhu, Sanja Fidler, and Xiaohui Zeng. Llama-mesh: Unifying 3d mesh generation with language models, 2024. URL <https://arxiv.org/abs/2411.09595>.
- Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, Chong Ruan, and Ping Luo. Janus: Decoupling visual encoding for unified multimodal understanding and generation. *arXiv preprint*, 2024.
- Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, Chong Ruan, et al. Janus: Decoupling visual encoding for unified multimodal understanding and generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 12966–12977, 2025a.
- Guanjun Wu, Jiemin Fang, Chen Yang, Sikuang Li, Taoran Yi, Jia Lu, Zanwei Zhou, Jiazhong Cen, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Xinggang Wang, and Qi Tian. Unilat3d: Geometry-appearance unified latents for single-stage 3d generation. *arXiv preprint arXiv:2509.25079*, 2025b.
- Shuang Wu, Youtian Lin, Feihu Zhang, Yifei Zeng, Yikang Yang, Jiachen Qian, Siyu Zhu, Xun Cao, Philip Torr, Yao Yao, et al. Direct3D-S2: Gigascale 3D generation made easy with spatial sparse attention. *Advances in Neural Information Processing Systems*, 38:170778–170804, 2026.
- Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation. *arXiv preprint arXiv:2412.01506*, 2024.
- Jianfeng Xiang, Xiaoxue Chen, Sicheng Xu, Ruicheng Wang, Zelong Lv, Yu Deng, Hongyuan Zhu, Yue Dong, Hao Zhao, Nicholas Jing Yuan, et al. Native and compact structured latents for 3d generation. *arXiv preprint arXiv:2512.14692*, 2025.
- Jinheng Xie, Weijia Mao, Zechen Bai, David Junhao Zhang, Weihao Wang, Kevin Qinghong Lin, Yuchao Gu, Zhiyu Chen, Zhenheng Yang, and Mike Zheng Shou. Show-o: One single transformer to unify multimodal understanding and generation. *arXiv preprint*, 2024.

- Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models, 2024a. URL <https://arxiv.org/abs/2404.07191>.
- Runsen Xu, Xiaolong Wang, Tai Wang, Yilun Chen, Jiangmiao Pang, and Dahua Lin. Pointllm: Empowering large language models to understand point clouds. In *European Conference on Computer Vision*, pages 131–147. Springer, 2024b.
- Le Xue, Mingfei Gao, Chen Xing, Roberto Martín-Martín, Jiajun Wu, Caiming Xiong, Ran Xu, Juan Carlos Niebles, and Silvio Savarese. Ulip: Learning a unified representation of language, images, and point clouds for 3d understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1179–1189, 2023.
- Le Xue, Ning Yu, Shu Zhang, Artemis Panagopoulou, Junnan Li, Roberto Martín-Martín, Jiajun Wu, Caiming Xiong, Ran Xu, Juan Carlos Niebles, et al. Ulip-2: Towards scalable multimodal pre-training for 3d understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27091–27101, 2024.
- Chongjie Ye, Yushuang Wu, Ziteng Lu, Jiahao Chang, Xiaoyang Guo, Jiaqing Zhou, Hao Zhao, and Xiaoguang Han. Hi3DGen: High-fidelity 3D geometry generation from images via normal bridging. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 25050–25061, 2025a.
- Chongjie Ye, Cheng Cao, Chuanyu Pan, Yiming Hao, Yihao Zhi, Yuanming Hu, and Xiaoguang Han. Omni123: Exploring 3d native foundation models with limited 3d data by unifying text to 2d and 3d generation. *arXiv preprint arXiv:2604.02289*, 2026.
- Junliang Ye, Zhengyi Wang, Ruowen Zhao, Shenghao Xie, and Jun Zhu. Shapellm-omni: A native multimodal llm for 3d generation and understanding. *arXiv preprint arXiv:2506.01853*, 2025b.
- Biao Zhang, Jiapeng Tang, Matthias Nießner, and Peter Wonka. 3dshape2vecset: A 3d shape representation for neural fields and generative diffusion models. *ACM Transactions on Graphics*, 42(4), July 2023. doi: 10.1145/3592442. URL <https://doi.org/10.1145/3592442>.
- Letian Zhang, Sucheng Ren, Yanqing Liu, Xianhang Li, Zeyu Wang, Yuyin Zhou, Huaxiu Yao, Zeyu Zheng, Weili Nie, Guilin Liu, et al. Openvision 3: A family of unified visual encoder for both understanding and generation. *arXiv preprint arXiv:2601.15369*, 2026.
- Yibo Zhang, Li Zhang, Rui Ma, and Nan Cao. Texverse: A universe of 3d objects with high-resolution textures, 2025. URL <https://arxiv.org/abs/2508.10868>.
- Zibo Zhao, Wen Liu, Xin Chen, Xianfang Zeng, Rui Wang, Pei Cheng, Bin Fu, Tao Chen, Gang Yu, and Shenghua Gao. Michelangelo: Conditional 3D shape generation based on shape-image-text aligned latent representation. *Advances in Neural Information Processing Systems*, 36:73969–73982, 2023.
- Junsheng Zhou, Jinsheng Wang, Baorui Ma, Yu-Shen Liu, Tiejun Huang, and Xinlong Wang. Uni3d: Exploring unified 3d representation at scale. In *International Conference on Learning Representations (ICLR)*, 2024.