

Novelty Assessment Report

Paper: FullPart: Generating each 3D Part at Full Resolution

PDF URL: <https://openreview.net/pdf?id=QIRIE7a1p4>

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-27

Abstract

Part-based 3D generation holds great potential for various applications. Previous part generators that represent parts using implicit vector-set tokens often suffer from insufficient geometric details. Another line of work adopts an explicit voxel representation but shares a global voxel grid among all parts; this often causes small parts to occupy too few voxels, leading to degraded quality. In this paper, we propose FullPart, a novel framework that combines both implicit and explicit paradigms. It first derives the bounding box layout through an implicit box vector-set diffusion process, a task that implicit diffusion handles effectively since box tokens contain little geometric detail. Then, it generates detailed parts, each within its own fixed full-resolution voxel grid. Instead of sharing a global low-resolution space, each part in our method—even small ones—is generated at full resolution, enabling the synthesis of intricate details. We further introduce a center-point encoding strategy to address the misalignment issue when exchanging information between parts of different actual sizes, thereby maintaining global coherence. Moreover, to tackle the scarcity of reliable part data, we present PartVerse-XL, the largest human-annotated 3D part dataset to date. Extensive experiments demonstrate that FullPart achieves state-of-the-art results in 3D part generation. We will release all code, data, and model to benefit future research in 3D part generation.

Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

Core Task Landscape

This paper addresses: **part-based 3D object generation with full-resolution representation**

A total of **33 papers** were analyzed and organized into a taxonomy with **20 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Part-Aware Representation and Decomposition Strategies**
- **Full-Resolution and High-Fidelity Volumetric Approaches**
- **Compositional and Multi-View Synthesis Methods**
- **Domain-Specific Part-Based Applications**
- **Representation Learning and Encoding Foundations**

Complete Taxonomy Tree

- part-based 3D object generation with full-resolution representation Survey Taxonomy
- Part-Aware Representation and Decomposition Strategies
 - Hierarchical and Multi-Scale Part Encoding
 - Octree and Sparse Hierarchical Representations (2 papers)
 - [2] Sparc3D: Sparse Representation and Construction for High-Resolution 3D Shapes Modeling (Li Zhi-Hao, 2025) [View paper](#)
 - [10] OctGPT: Octree-based Multiscale Autoregressive Models for 3D Shape Generation (Zhou Chuan-zhi, 2025) [View paper](#)
 - Gaussian Mixture and Hierarchical Gaussian Representations (2 papers)
 - [4] Ultra3D: Efficient and High-Fidelity 3D Generation with Part Attention (Chen Yi-Wen, 2025) [View paper](#)
 - [23] Hierarchical Gaussian Mixture Model Splatting for Efficient and Part Controllable 3D Generation (Qitong Yang, 2025) [View paper](#)
 - Part-Level Latent Diffusion and Generative Models
 - Contextual Part Latent Diffusion ★ (2 papers)
 - [0] FullPart: Generating each 3D Part at Full Resolution (Anon et al., 2026) [View paper](#)
 - [3] From one to more: Contextual part latents for 3d generation (Dong, 2025) [View paper](#)
 - Part Sampling and Fragment-Based Generation (2 papers)
 - [8] Chirpy3D: Creative Fine-grained 3D Object Fabrication via Part Sampling (Ng, 2025) [View paper](#)
 - [11] Diverse Part Synthesis for 3D Shape Creation (Yanran Guan, 2024) [View paper](#)
 - Variational Part Autoencoders (1 papers)
 - [14] The shape variational autoencoder: A deep generative model of part-aware segmented 3d objects (C Nash, 2017) [View paper](#)
 - Part Assembly and Spatial Arrangement
 - Anchor-Based and Graph-Structured Assembly (2 papers)
 - [7] Assembler: Scalable 3D Part Assembly via Anchor Point Diffusion (Zhao Wang, 2025) [View paper](#)
 - [16] SceneHGN: Hierarchical Graph Networks for 3D Indoor Scene Generation With Fine-Grained Geometry (Lin Gao, 2023) [View paper](#)
 - Part Retrieval and Structure Recovery (2 papers)
 - [30] Structure recovery by part assembly (Chao-Hui Shen, 2012) [View paper](#)
 - [33] 3D Synthesis of Man-made Objects based on Fine-grained Parts (Becerra, 2018) [View paper](#)
- Full-Resolution and High-Fidelity Volumetric Approaches
 - Dual Volume and Component-Aware Packing (2 papers)

- [5] Efficient Part-level 3D Object Generation via Dual Volume Packing (Tang Jia-xiang, 2025) [View paper](#)
- [31] Component-Aware High-Resolution 3D Object Reconstruction (Weichao Shen, 2021) [View paper](#)
- Nested and Matryoshka Volumetric Encodings (1 papers)
- [26] Matryoshka Networks: Predicting 3D Geometry via Nested Shape Layers (Stephan R. Richter, 2018) [View paper](#)
- Diffusion Time-Step and Hash-Encoding Curricula (1 papers)
- [6] Diffusion time-step curriculum for one image to 3d generation (Yi, 2024) [View paper](#)
- Compositional and Multi-View Synthesis Methods
 - Category-Aware and Multi-Modal Composition (2 papers)
 - [1] Category-Aware 3D Object Composition with Disentangled Texture and Shape Multi-view Diffusion (Chen, 2025) [View paper](#)
 - [24] Automated 3D Reconstruction of Mechanical Components Using Multimodal Deep Learning (Issam Dridi, 2025) [View paper](#)
 - Multi-View Diffusion and Amodal Reconstruction (2 papers)
 - [18] AmodalGen3D: Generative Amodal 3D Object Reconstruction from Sparse Unposed Views (Junwei Zhou, 2025) [View paper](#)
 - [21] 3D-assisted feature synthesis for novel views of an object (Hao Su, 2015) [View paper](#)
 - Generative Scene and Object Decomposition (2 papers)
 - [17] Deep Generative Modeling Based on VAE-GAN for 3D Indoor Scene Synthesis (Shuai Li, 2023) [View paper](#)
 - [22] Generative 4D Scene Gaussian Splatting with Object View-Synthesis Priors (Ke Lei, 2025) [View paper](#)
- Domain-Specific Part-Based Applications
 - Human Body and Articulated Shape Modeling (2 papers)
 - [20] GHUM & GHUML: Generative 3D Human Shape and Articulated Pose Models (Hongyi Xu, 2020) [View paper](#)
 - [28] MonoCloth: Reconstruction and Animation of Cloth-Decoupled Human Avatars from Monocular Videos (Ying, 2025) [View paper](#)
 - Facial and Portrait Part-Based Synthesis
 - 3D-Aware Compositional Portrait Generation (2 papers)
 - [9] 3D-SSGAN: Lifting 2D Semantics for 3D-Aware Compositional Portrait Synthesis (Liu Ruiqi, 2024) [View paper](#)
 - [15] Semantic 3D-aware Portrait Synthesis and Manipulation Based on Compositional Neural Radiance Field (Dong Jing, 2023) [View paper](#)
 - Deformation-Based Talking Head Synthesis (2 papers)
 - [13] TalkingGaussian: Structure-Persistent 3D Talking Head Synthesis via Gaussian Splatting (Jiahe Li, 2024) [View paper](#)
 - [29] PoseTalker: Talking Head Synthesis via Pose-Driven 3D Gaussian Splatting (Zhen Xiong, 2025) [View paper](#)
 - Part-Based Local-Global Facial Transfer (1 papers)
 - [12] High-fidelity facial expression transfer using part-based local-global conditional gans (Muhammad Mamunur Rashid, 2023) [View paper](#)
- Representation Learning and Encoding Foundations
 - Discrete Representation and Hash Encoding (1 papers)
 - [25] Discrete Representation Learning for Retrieval and Generation (Ng, 2025) [View paper](#)
 - Fragment and Part-Based Object Representation (2 papers)
 - [19] A fragment-based approach to object representation and classification (Shimon Ullman, 2001) [View paper](#)
 - [32] Completions: Reuse and Object Representations (Daniel Norell, 2021) [View paper](#)
 - Generative Shape and Texture Modeling (1 papers)
 - [27] Synthesis of 3D MRI Brain Images With Shape and Texture Generative Adversarial Deep Neural Networks (Chee Keong-Chong, 2021) [View paper](#)

Narrative

Core task: part-based 3D object generation with full-resolution representation. The field organizes around several complementary directions. Part-Aware Representation and Decomposition Strategies focus on how to segment and model objects as assemblies of meaningful components, often leveraging latent diffusion or hierarchical encodings to capture part-level semantics. Full-Resolution and High-Fidelity Volumetric Approaches emphasize maintaining geometric detail at scale, exploring techniques such as dual volume packing (Dual Volume Packing[5]) or octree-based representations (OctGPT[10]) to balance memory and fidelity. Compositional and Multi-View Synthesis Methods address the challenge of generating coherent 3D content from multiple viewpoints or by composing learned part priors, with works like Category Aware Composition[1] and Sparc3D[2] illustrating how category-specific knowledge can guide assembly. Domain-Specific Part-Based Applications tailor these ideas to specialized contexts—ranging from human body modeling (GHUM[20]) and talking avatars (TalkingGaussian[13], PoseTalker[29]) to medical imaging (3D MRI Synthesis[27])—while Representation Learning and Encoding Foundations provide the underlying machinery, including variational autoencoders (Shape VAE[14]) and discrete tokenization schemes (Discrete Representation Learning[25]).

A particularly active line of work explores part-level latent diffusion, where generative models operate in a structured latent space that respects object decomposition. Within this cluster, Contextual Part Latents[3] conditions diffusion on part relationships to ensure coherent assembly, while FullPart[0] extends this idea by maintaining full-resolution detail throughout the generation process, avoiding the loss of fine geometric features that can occur with coarser representations. Nearby efforts such as Diverse Part Synthesis[11] and Assembler[7] similarly emphasize compositional generation but may trade off resolution for broader part diversity or faster sampling. The central tension across these branches is between expressive part-level control and the computational cost of high-fidelity volumetric outputs. FullPart[0] sits at the intersection of contextual part modeling and full-resolution synthesis, aiming to preserve both semantic decomposition and geometric detail—a balance that distinguishes it from methods prioritizing either coarse part assembly or resolution alone.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. From one to more: Contextual part latents for 3d generation

Authors: Dong, Shaocong, Ding, Lihe, Chen Xiao, et al. (18 authors total) | **Year/Venue:** 2025 | **URL:** [View paper](#)

Abstract

Recent advances in 3D generation have transitioned from multi-view 2D rendering approaches to 3D-native latent diffusion frameworks that exploit geometric priors in ground truth data. Despite progress, three key limitations persist: (1) Single-latent representations fail to capture complex multi-part geometries, causing detail degradation; (2) Holistic latent coding neglects part independence and interrelationships critical for compositional design; (3) Global conditioning mechanisms lack fine-g...

Relationship Analysis

Both papers belong to the Contextual Part Latent Diffusion category, using latent representations to model inter-part relationships in 3D generation. They overlap in representing objects as multiple part latents and employing diffusion frameworks for coherent part generation. However, the original paper (FullPart) generates each part at full resolution in isolated voxel grids with center-corner encoding to handle scale misalignment, while the candidate paper (CoPart) uses hybrid geometric and image tokens with mutual guidance between modalities and parts, without the isolated full-resolution grid strategy.

Contributions Analysis

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: FullPart framework combining implicit and explicit paradigms

Description: The authors introduce FullPart, a framework that first generates bounding box layouts using implicit vecset diffusion, then generates each part at full resolution within its own dedicated voxel grid using explicit representation. This design addresses limitations of prior methods by enabling fine geometric details while maintaining global coherence.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Implicit-Explicit Coupling Enhancement for UAV Scene 3D Reconstruction

URL: [View paper](#)

Brief Assessment

Implicit-Explicit Coupling[56] focuses on UAV scene reconstruction using implicit-explicit coupling for photometric consistency and point cloud density enhancement, not part-level 3D object generation with layout-to-part synthesis.

2. PARIS: Part-level Reconstruction and Motion Analysis for Articulated Objects

URL: [View paper](#)

Brief Assessment

PARIS[58] focuses on articulated object reconstruction from multi-view images with motion analysis, not on part-level 3D generation combining implicit layout with explicit part structure generation as in the original paper.

3. LEIA: Latent View-invariant Embeddings for Implicit 3D Articulation

URL: [View paper](#)

Brief Assessment

LEIA[60] focuses on dynamic 3D object articulation using hypernetworks and view-invariant latent representations, not on part-level generation combining implicit layout with explicit structure generation.

4. Generating Part-Aware Editable 3D Shapes without 3D Supervision

URL: [View paper](#)

Brief Assessment

Part-Aware Editable[54] focuses on part-aware NeRF generation without 3D supervision, using implicit neural representations throughout. It does not combine implicit layout with explicit voxel-based part generation as FullPart does.

5. Learning Explicit Contact for Implicit Reconstruction of Hand-held Objects from Monocular Images

URL: [View paper](#)

Brief Assessment

Explicit Contact Learning[55] focuses on hand-held object reconstruction from monocular images using contact prediction between hands and objects, not on part-level 3D generation with layout-to-structure pipelines. The technical approaches and problem domains are fundamentally different.

6. SENS: Part-Aware Sketch-based Implicit Neural Shape Modeling

URL: [View paper](#)

Brief Assessment

SENS[59] focuses on sketch-based 3D shape generation using part-aware neural implicit representations, not on combining implicit layout diffusion with explicit voxel-based part generation as in the original paper.

7. Implicit Neural Head Synthesis via Controllable Local Deformation Fields

URL: [View paper](#)

Brief Assessment

Controllable Local Deformation[57] focuses on neural head synthesis using local deformation fields for facial animation, not on part-level 3D object generation with layout-to-part pipelines.

8. Neural parts: Learning expressive 3d shape abstractions with invertible neural networks

URL: [View paper](#)

Brief Assessment

Neural Parts[52] focuses on learning homeomorphic mappings between spheres and target objects to create expressive 3D primitives for shape abstraction. It does not address part-level 3D generation with layout-then-structure pipelines or the combination of implicit layout representation with explicit part generation at full resolution, which are the core innovations of the original paper.

9. Partsdf: Part-based implicit neural representation for composite 3d shape parametrization and optimization

URL: [View paper](#)

Brief Assessment

Partsdf[51] focuses on supervised implicit neural representations for composite 3D shapes using signed distance functions, not on combining implicit layout generation with explicit part structure generation as in the original paper's diffusion-based approach.

10. Anise: Assembly-based neural implicit surface reconstruction

URL: [View paper](#)

Brief Assessment

Anise[53] focuses on part-aware reconstruction from partial observations using neural implicit functions for assembly, not on combining implicit layout generation with explicit voxel-based part generation for 3D synthesis as in FullPart.

Contribution 2: Center-corner encoding strategy for part coherence

Description: The authors propose a center-corner encoding mechanism that embeds absolute spatial context for each voxel by encoding the positions of its center and eight corners in a unified super-high-resolution global coordinate system. This addresses the scale misalignment problem when parts of different sizes exchange information through attention mechanisms.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. From one to more: Contextual part latents for 3d generation

URL: [View paper](#)

Brief Assessment

Contextual Part Latents[3] uses mutual guidance with cross-part and cross-modality attention mechanisms to ensure part coherence, rather than center-corner positional encoding. The technical approaches for maintaining global coherence differ fundamentally.

2. PartCrafter: Structured 3D Mesh Generation via Compositional Latent Diffusion Transformers

URL: [View paper](#)

Brief Assessment

PartCrafter[45] uses part identity embeddings and local-global attention mechanisms for part coherence, but does not employ a center-corner encoding strategy that embeds absolute spatial context through corner positions in a super-high-resolution global coordinate system as described in the original paper.

3. Romantex: Decoupling 3d-aware rotary positional embedded multi-attention network for texture synthesis

URL: [View paper](#)

Brief Assessment

Romantex[43] focuses on texture synthesis for 3D assets using rotary positional embeddings in a multi-view diffusion framework, not on part-based generation with spatial encoding for maintaining coherence across parts of different sizes.

4. Videograin: Modulating space-time attention for multi-grained video editing

URL: [View paper](#)

Brief Assessment

Videograin[46] addresses video editing with spatial-temporal attention modulation for multi-grained control across video frames, not 3D part generation with voxel grids and spatial encoding for maintaining coherence across parts of different sizes.

5. Enhanced Monocular Depth Estimation Based on Improved Self-Attention Mechanisms and Composite Loss Functions

URL: [View paper](#)

Brief Assessment

Enhanced Monocular Depth[48] focuses on depth estimation using self-attention mechanisms for 2D image processing, not 3D part generation with spatial encoding for voxel-based representations.

6. mpAuvS: multi-perspective attention for unsupervised video summarization capturing global, local, and spatiotemporal context: C. Xin et al.

URL: [View paper](#)

Brief Assessment

mpAuvS[49] focuses on video summarization with attention mechanisms for temporal content, not 3D part generation with spatial encoding for maintaining coherence across parts of different sizes.

7. Omnipart: Part-aware 3d generation with semantic decoupling and structural cohesion

URL: [View paper](#)

Brief Assessment

Omnipart[44] uses part position embeddings (PPE) and bounding box-based spatial conditioning, but does not employ a center-corner encoding mechanism that embeds absolute spatial context by encoding center and eight corner positions in a super-high-resolution global coordinate system as described in the original paper.

8. Multi-Modality Regional Alignment Network for Covid X-Ray Survival Prediction and Report Generation

URL: [View paper](#)

Brief Assessment

Multi Modality Regional[47] focuses on medical image analysis (COVID X-ray survival prediction and report generation) using regional alignment for radiology tasks, not 3D part generation with spatial encoding for maintaining global coherence in attention mechanisms.

9. Ultra3D: Efficient and High-Fidelity 3D Generation with Part Attention

URL: [View paper](#)

Brief Assessment

Ultra3D[4] focuses on part attention mechanisms for computational efficiency in sparse voxel generation, not on spatial encoding strategies for maintaining scale alignment between parts of different sizes.

10. Directional Non-Commutative Monoidal Structures for Compositional Embeddings in Machine Learning

URL: [View paper](#)

Brief Assessment

Directional Monoidal Structures[50] addresses compositional embeddings through algebraic operators for multi-dimensional data, not spatial encoding for part-based 3D generation. The candidate focuses on abstract mathematical frameworks for sequence/grid modeling, while the original addresses concrete voxel-level spatial alignment in 3D part generation.

Contribution 3: PartVerse-XL dataset

Description: The authors introduce PartVerse-XL, the largest human-annotated 3D part dataset to date, containing 40K objects and 320K parts with associated part-aware texture descriptions. The dataset was created through mesh pre-segmentation followed by human refinement to ensure high-quality, semantically consistent annotations.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. From one to more: Contextual part latents for 3d generation

URL: [View paper](#)

Prior Art Analysis

Contextual Part Latents[3] introduces PartVerse, a large-scale 3D part dataset from Objaverse with human post-annotations. The paper describes collecting 91k parts from 12k objects across 175 categories using automatic mesh segmentation followed by human refinement - a methodology nearly identical to the original paper's PartVerse-XL creation process. Both datasets use Objaverse as source, employ automatic pre-segmentation with human post-processing, and emphasize semantic consistency through manual annotation. The scale differs (PartVerse: 91k parts/12k objects vs. PartVerse-XL: 320k parts/40k objects), but the fundamental contribution of creating a large human-annotated part dataset with this methodology was already demonstrated.

Evidence

Evidence 1 - **Rationale:** Both papers claim to introduce large-scale human-annotated part datasets from Objaverse, with Contextual Part Latents[3] published earlier demonstrating this contribution type. - **Original:** we present partverse-xl, the largest human-annotated 3d part dataset to date with 40k objects and 320k parts - **Candidate:** we introduce partverse, a new diverse 3d part dataset comprising 91k high-quality parts from 12k objects with detailed text descriptions

Evidence 2 - **Rationale:** Both datasets use Objaverse as the source and employ similar curation methodologies with human annotation for quality. - **Original:** we selected 40k objects from objaverse-xl (deitke et al., 2023), and create a high-quality part annotation, using mesh pre-segmentation followed by human refinement to ensure high-quality, semantically consistent annotations - **Candidate:** our partverse, curated and annotated from objaverse [7], exhibits enhanced diversity in 175 categories and realistic textures, significantly improving the model's ability to generate high-quality 3d parts

2. PartNeXt: A Next-Generation Dataset for Fine-Grained and Hierarchical 3D Part Understanding

URL: [View paper](#)

Brief Assessment

PartNeXt[39] focuses on a different dataset (23K textured models across 50 categories) with different annotation methodology and benchmarking tasks (class-agnostic segmentation and 3D-LLM question answering), rather than the 40K objects with 320K parts and part-aware texture descriptions in PartVerse-XL.

3. MinD-3D++: Advancing fMRI-Based 3D Reconstruction With High-Quality Textured Mesh Generation and a Comprehensive Dataset

URL: [View paper](#)

Brief Assessment

MinD-3D[35] introduces fMRI-3D dataset for brain-to-3D reconstruction tasks, not a human-annotated 3D part dataset with part-aware texture descriptions for generative modeling.

4. A large-scale annotated mechanical components benchmark for classification and retrieval tasks with deep neural networks

URL: [View paper](#)

Brief Assessment

Mechanical Components Benchmark[34] focuses on mechanical engineering components (bearings, gears, fasteners) for industrial classification, while PartVerse-XL addresses general 3D object part segmentation with semantic descriptions. The domains, annotation methodologies, and intended applications are fundamentally different.

5. Multimodal Large Language Models: A Survey

URL: [View paper](#)

Brief Assessment

Multimodal LLM Survey[42] is a broad survey paper covering generative models across multiple modalities (text, image, music, video, motion, 3D). While it discusses 3D object generation techniques, it does not present a specific human-annotated 3D part dataset comparable to PartVerse-XL.

6. Clevrer-humans: Describing physical and causal events the human way

URL: [View paper](#)

Brief Assessment

Clevrer Humans[37] focuses on human-annotated causal event descriptions in physical reasoning videos, not 3D part segmentation datasets. The domains are fundamentally different (video reasoning vs. 3D object part annotation).

7. Segment Any 3D-Part in a Scene from a Sentence

URL: [View paper](#)

Brief Assessment

Segment Any 3D[41] introduces the 3D-pu dataset for 3D scene understanding with part annotations, while PartVerse-XL focuses on individual 3D object part generation. The datasets serve different purposes: 3D-pu is designed for scene-level part segmentation tasks, whereas PartVerse-XL supports part-based 3D object generation with texture descriptions.

8. CityRefer: Geography-aware 3D Visual Grounding Dataset on City-scale Point Cloud Data

URL: [View paper](#)

Brief Assessment

CityRefer[38] focuses on city-scale 3D visual grounding with outdoor scene annotations, not part-level 3D object datasets with part-aware texture descriptions.

9. PARIS3D: Reasoning-based 3D Part Segmentation Using Large Multimodal Model

URL: [View paper](#)

Brief Assessment

PARIS3D[36] focuses on reasoning-based segmentation with textual queries and presents a dataset of 60k instructions for reasoning tasks, not a large-scale human-annotated 3D part dataset with part-aware texture descriptions like PartVerse-XL.

10. Affogato: Learning Open-Vocabulary Affordance Grounding with Automated Data Generation at Scale

URL: [View paper](#)

Brief Assessment

Affogato[40] focuses on affordance grounding with interaction-based annotations, not part segmentation datasets. The datasets serve fundamentally different purposes: PartVerse-XL provides part decomposition with texture descriptions, while Affogato provides affordance heatmaps for human-object interactions.

Appendix: Text Similarity Detection

Textual similarity detection checked 29 papers and found 1 similarity segment(s) across 1 paper(s).

The following **1 paper(s)** were detected to have high textual similarity with the original paper. These may represent different versions of the same work, duplicate submissions, or papers with substantial textual overlap. Readers are advised to verify these relationships independently.

1. Omnipart: Part-aware 3d generation with semantic decoupling and structural cohesion

Detected in: Contribution: contribution_2

△ **Note:** This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

References

- [0] FullPart: Generating each 3D Part at Full Resolution [View paper](#)
- [1] Category-Aware 3D Object Composition with Disentangled Texture and Shape Multi-view Diffusion [View paper](#)
- [2] Sparc3D: Sparse Representation and Construction for High-Resolution 3D Shapes Modeling [View paper](#)
- [3] From one to more: Contextual part latents for 3d generation [View paper](#)
- [4] Ultra3D: Efficient and High-Fidelity 3D Generation with Part Attention [View paper](#)
- [5] Efficient Part-level 3D Object Generation via Dual Volume Packing [View paper](#)
- [6] Diffusion time-step curriculum for one image to 3d generation [View paper](#)
- [7] Assembler: Scalable 3D Part Assembly via Anchor Point Diffusion [View paper](#)
- [8] Chirpy3D: Creative Fine-grained 3D Object Fabrication via Part Sampling [View paper](#)
- [9] 3D-SSGAN: Lifting 2D Semantics for 3D-Aware Compositional Portrait Synthesis [View paper](#)
- [10] OctGPT: Octree-based Multiscale Autoregressive Models for 3D Shape Generation [View paper](#)
- [11] Diverse Part Synthesis for 3D Shape Creation [View paper](#)
- [12] High-fidelity facial expression transfer using part-based local+global conditional gans [View paper](#)
- [13] TalkingGaussian: Structure-Persistent 3D Talking Head Synthesis via Gaussian Splatting [View paper](#)
- [14] The shape variational autoencoder: A deep generative model of part+segmented 3d objects [View paper](#)
- [15] Semantic 3D-aware Portrait Synthesis and Manipulation Based on Compositional Neural Radiance Field [View paper](#)
- [16] SceneHGN: Hierarchical Graph Networks for 3D Indoor Scene Generation With Fine-Grained Geometry [View paper](#)
- [17] Deep Generative Modeling Based on VAE-GAN for 3D Indoor Scene Synthesis [View paper](#)
- [18] AmodalGen3D: Generative Amodal 3D Object Reconstruction from Sparse Unposed Views [View paper](#)
- [19] A fragment-based approach to object representation and classification [View paper](#)
- [20] GHUM & GHUML: Generative 3D Human Shape and Articulated Pose Models [View paper](#)
- [21] 3D-assisted feature synthesis for novel views of an object [View paper](#)
- [22] Generative 4D Scene Gaussian Splatting with Object View-Synthesis Priors [View paper](#)
- [23] Hierarchical Gaussian Mixture Model Splatting for Efficient and Part Controllable 3D Generation [View paper](#)
- [24] Automated 3D Reconstruction of Mechanical Components Using Multimodal Deep Learning [View paper](#)
- [25] Discrete Representation Learning for Retrieval and Generation [View paper](#)
- [26] Matryoshka Networks: Predicting 3D Geometry via Nested Shape Layers [View paper](#)
- [27] Synthesis of 3D MRI Brain Images With Shape and Texture Generative Adversarial Deep Neural Networks [View paper](#)
- [28] MonoCloth: Reconstruction and Animation of Cloth-Decoupled Human Avatars from Monocular Videos [View paper](#)
- [29] PoseTalker: Talking Head Synthesis via Pose-Driven 3D Gaussian Splatting [View paper](#)
- [30] Structure recovery by part assembly [View paper](#)
- [31] Component-Aware High-Resolution 3D Object Reconstruction [View paper](#)
- [32] Completions: Reuse and Object Representations [View paper](#)
- [33] 3D Synthesis of Man-made Objects based on Fine-grained Parts [View paper](#)
- [34] A large-scale annotated mechanical components benchmark for classification and retrieval tasks with deep neural networks [View paper](#)
- [35] MinD-3D++: Advancing fMRI-Based 3D Reconstruction With High-Quality Textured Mesh Generation and a Comprehensive Dataset [View paper](#)
- [36] PARIS3D: Reasoning-based 3D Part Segmentation Using Large Multimodal Model [View paper](#)
- [37] Clevrer-humans: Describing physical and causal events the human way [View paper](#)
- [38] CityRefer: Geography-aware 3D Visual Grounding Dataset on City-scale Point Cloud Data [View paper](#)
- [39] PartNeXt: A Next-Generation Dataset for Fine-Grained and Hierarchical 3D Part Understanding [View paper](#)
- [40] Affogato: Learning Open-Vocabulary Affordance Grounding with Automated Data Generation at Scale [View paper](#)
- [41] Segment Any 3D-Part in a Scene from a Sentence [View paper](#)
- [42] Multimodal Large Language Models: A Survey [View paper](#)
- [43] Romantex: Decoupling 3d-aware rotary positional embedded multi-attention network for texture synthesis [View paper](#)
- [44] Omnipart: Part-aware 3d generation with semantic decoupling and structural cohesion [View paper](#)
- [45] PartCrafter: Structured 3D Mesh Generation via Compositional Latent Diffusion Transformers [View paper](#)
- [46] Videograin: Modulating space-time attention for multi-grained video editing [View paper](#)
- [47] Multi-Modality Regional Alignment Network for Covid X-Ray Survival Prediction and Report Generation [View paper](#)

- [48] Enhanced Monocular Depth Estimation Based on Improved Self-Attention Mechanisms and Composite Loss Functions [View paper](#)
- [49] mpAuvS: multi-perspective attention for unsupervised video summarization capturing global, local, and spatiotemporal context: C. Xin et al. [View paper](#)
- [50] Directional Non-Commutative Monoidal Structures for Compositional Embeddings in Machine Learning [View paper](#)
- [51] PartsdF: Part-based implicit neural representation for composite 3d shape parametrization and optimization [View paper](#)
- [52] Neural parts: Learning expressive 3d shape abstractions with invertible neural networks [View paper](#)
- [53] Anise: Assembly-based neural implicit surface reconstruction [View paper](#)
- [54] Generating Part-Aware Editable 3D Shapes without 3D Supervision [View paper](#)
- [55] Learning Explicit Contact for Implicit Reconstruction of Hand-held Objects from Monocular Images [View paper](#)
- [56] Implicit-Explicit Coupling Enhancement for UAV Scene 3D Reconstruction [View paper](#)
- [57] Implicit Neural Head Synthesis via Controllable Local Deformation Fields [View paper](#)
- [58] PARIS: Part-level Reconstruction and Motion Analysis for Articulated Objects [View paper](#)
- [59] SENS: Part-Aware Sketch-based Implicit Neural Shape Modeling [View paper](#)
- [60] LEIA: Latent View-invariant Embeddings for Implicit 3D Articulation [View paper](#)