Novelty Assessment Report

Paper: HoloPart: Generative 3D Part Amodal Segmentation
PDF URL: https://openreview.net/pdf?id=2VsBJwefDC
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2026-01-05

Abstract

3D part amodal segmentation--decomposing a 3D shape into complete, semantically meaningful parts, even when occluded--is a challenging but crucial task for 3D content creation and understanding. Existing 3D part segmentation methods only identify visible surface patches, limiting their utility. Inspired by 2D amodal segmentation, we introduce this novel task to the 3D domain and propose a practical, two-stage approach, addressing the key challenges of inferring occluded 3D geometry, maintaining global shape consistency, and handling diverse shapes with limited training data. First, we leverage existing 3D part segmentation to obtain initial, incomplete part segments. Second, we introduce HoloPart, a novel diffusion-based model, to complete these segments into full 3D parts. HoloPart utilizes a specialized architecture with local attention to capture fine-grained part geometry and global shape context attention to ensure overall shape consistency. We introduce new benchmarks based on the ABO and PartObjaverse-Tiny datasets and demonstrate that HoloPart significantly outperforms state-of-the-art shape completion methods. By incorporating HoloPart with existing segmentation techniques, we achieve promising results on 3D part amodal segmentation, opening new avenues for applications in geometry editing, animation, and material assignment.

Core Task Landscape

This paper addresses: **3D Part Amodal Segmentation with Complete Geometry Inference**

A total of **50 papers** were analyzed and organized into a taxonomy with **20 categories**.

Taxonomy Overview

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

- **Amodal Completion and Shape Reconstruction Methods**
- **Part-Level Segmentation and Decomposition**
- **2D Amodal Segmentation and Image-Based Methods**
- **Instance Segmentation and Tracking with Occlusion Handling**
- **Geometry-Guided Segmentation and Reconstruction**
- **Domain-Specific Applications**

Complete Taxonomy Tree

- 3D Part Amodal Segmentation with Complete Geometry Inference Survey Taxonomy
- Amodal Completion and Shape Reconstruction Methods
  - Diffusion-Based Amodal Reconstruction ★ (5 papers)
  - [0] HoloPart: Generative 3D Part Amodal Segmentation (Anon et al., 2026) View paper
  - [1] pix2gestalt: Amodal segmentation by synthesizing wholes (Ege Ozguroglu, 2024) View paper
  - [8] Amodal3r: Amodal 3d reconstruction from occluded 2d images (Wu, 2025) View paper
  - [10] AmodalGen3D: Generative Amodal 3D Object Reconstruction from Sparse Unposed Views (Junwei Zhou, 2025) View paper
  - [27] Sequential Amodal Segmentation via Cumulative Occlusion Learning (Ao, 2024) View paper
  - Other Generative Amodal Approaches (2 papers)
  - [20] Variational amodal object completion (Huan Ling, 2020) View paper
  - [29] GIN: Generative invariant shape prior for amodal instance segmentation (Zhixuan Li, 2023) View paper
  - Multi-View and Video-Based Amodal Reconstruction (3 papers)
  - [13] Viser: Video-specific surface embeddings for articulated 3d shape reconstruction (Gengshan Yang, 2021) View paper
  - [33] Self-supervised Amodal Video Object Segmentation (Yao Jian, 2022) View paper
  - [48] A2VIS: Amodal-Aware Approach to Video Instance Segmentation (Minh Tran, 2024) View paper
  - Single-View Amodal Reconstruction (2 papers)
  - [21] Monogrnet: A geometric reasoning network for monocular 3d object localization (Zengyi Qin, 2019) View paper
  - [39] Amodal detection of 3d objects: Inferring 3d bounding boxes from 2d ones in rgb-depth images (Zhuo Deng, 2017) View paper
- Part-Level Segmentation and Decomposition
  - Part Segmentation with Amodal Reasoning (4 papers)
  - [4] PartSAM: A Scalable Promptable Part Segmentation Model Trained on Native 3D Data (Zhu Zhe, 2025) View paper
  - [5] Stable Part Diffusion 4D: Multi-View RGB and Kinematic Parts Video Generation (Zhang Hao, 2025) View paper
  - [6] Omnipart: Part-aware 3d generation with semantic decoupling and structural cohesion (Yang, 2025) View paper
  - [7] PartGen: Part-level 3D Generation and Reconstruction with Multi-View Diffusion Models (Minghao Chen, 2024) View paper
  - Part Segmentation without Amodal Reasoning (3 papers)
  - [2] Large-Scale 3D Shape Reconstruction and Segmentation from ShapeNet Core55 (Yi Li, 2022) View paper
  - [30] ASIA: Adaptive 3D Segmentation using Few Image Annotations (Perla Sai Raj Kishore, 2025) View paper
  - [40] 3D Part Segmentation via Geometric Aggregation of 2D Visual Features (Marco Garosi, 2024) View paper

    - Articulated Object Part Decomposition (3 papers)
    - [12] DreamArt: Generating Interactable Articulated Objects from a Single Image (Lu RuiJie, 2025) View paper
    - [36] KineDiff3D: Kinematic-Aware Diffusion for Category-Level Articulated Object Shape Reconstruction and Generation (Xu Wenbo, 2025) View paper
    - [46] Generating Objects with Part-Articulation from a Single Image (Ruijie Lu, 2025) View paper
    - Part-Aware 3D Generation (2 papers)
    - [16] BANG: Dividing 3D Assets via Generative Exploded Dynamics (Zhang Long-wen, 2025) View paper
    - [18] Imagine: Image-Guided 3D Part Assembly with Structure Knowledge Graph (He Bin, 2025) View paper
- 2D Amodal Segmentation and Image-Based Methods
    - Foundation Model-Based 2D Amodal Segmentation (3 papers)
    - [24] PLUG: Revisiting Amodal Segmentation with Foundation Model and Hierarchical Focus (Liu Zhao-chen, 2024) View paper
    - [44] Single Point, Full Mask: Velocity-Guided Level Set Evolution for End-to-End Amodal Segmentation (Li Zhixuan, 2025) View paper
    - [45] ViTA-Seg: Vision Transformer for Amodal Segmentation in Robotics (Donato Caramia, 2025) View paper
    - Task-Specific 2D Amodal Segmentation (3 papers)
    - [11] Application of amodal segmentation for shape reconstruction and occlusion recovery in occluded tomatoes (Jing Yang, 2024) View paper
    - [23] Application of amodal segmentation for shape (H Deng, 2025) View paper
    - [26] Learning Vector Quantized Shape Code for Amodal Blastomere Instance Segmentation (Jang, 2023) View paper
    - Semantic and Panoptic 2D Amodal Segmentation (2 papers)
    - [19] Semantic Amodal Segmentation (Yan Zhu, 2022) View paper
    - [34] Perceiving the Invisible: Proposal-Free Amodal Panoptic Segmentation (Rohit Mohan, 2022) View paper
    - Weakly-Supervised 2D Amodal Segmentation (1 papers)
    - [38] BLADE: Box-Level Supervised Amodal Segmentation through Directed Expansion (Jiang Tingting, 2024) View paper
- Instance Segmentation and Tracking with Occlusion Handling (3 papers)
    - [3] 3D-Aware Instance Segmentation and Tracking in Egocentric Videos (Yash Bhalgat, 2024) View paper
    - [31] NeuralDiff: Segmenting 3D objects that move in egocentric videos (Tschernezki, 2021) View paper
    - [50] Live Reconstruction of Large-Scale Dynamic Outdoor Worlds (Miksik, 2019) View paper
- Geometry-Guided Segmentation and Reconstruction
    - Geometric Structure Modeling for Segmentation (3 papers)
    - [17] Unsupervised 3D Object Segmentation of Point Clouds by Geometry Consistency (Ziyang Song, 2024) View paper
    - [25] GEOMetrics: Exploiting Geometric Structure for Graph-Encoded Objects (Smith, 2022) View paper
    - [35] Holistic Pose Graph: Modeling Geometric Structure among Objects in a Scene using Graph Inference for 3D Object Prediction (Jiwei Xiao, 2021) View paper
    - Shape Prior-Based Reconstruction (2 papers)
    - [14] 3D Face Reconstruction with the Geometric Guidance of Facial Part Segmentation (Zidu Wang, 2023) View paper
    - [32] Part2Point: A Part-Oriented Point Cloud Reconstruction Framework (Yu-Cheng Feng, 2023) View paper
- Domain-Specific Applications
    - Agricultural and Botanical Applications (4 papers)
    - [9] Obscured tree branches segmentation and 3D reconstruction using deep learning and geometrical constraints (Eugene Kok, 2023) View paper
    - [22] Multi-vision-based Localization and Pose Estimation of Occluded Apple Fruits for Harvesting Robots (Tao Li, 2022) View paper
    - [28] Development of a Precise Tree Structure from LiDAR Point Clouds (Abdul Nurunnabi, 2024) View paper
    - [49] ENHANCED DETECTION OF 3D INDIVIDUAL TREES IN FORESTED AREAS USING AIRBORNE FULL-WAVEFORM LIDAR DATA BY COMBINING NORMALIZED CUTS WITH SPATIAL DENSITY CLUSTERING (Wei Yao, 2013) View paper
    - Robotic Manipulation and Grasping (2 papers)
    - [42] ImVoxelGNet: Image to voxels geometry-aware projection for multi-view RGB-based 3D object detection. (Gang Xu, 2025) View paper
    - [43] 3D deep object recognition and semantic understanding for visually-guided robotic service (Sukhan Lee, 2018) View paper
    - Medical and Biological Imaging (2 papers)
    - [37] Poster 140: Patella Shape is Associated with ACL Injury and Changes in KOOS Following ACL Reconstruction (James Peters, 2025) View paper
    - [41] Structure-Guided Segmentation for 3D Neuron Reconstruction. (Bo Yang, 2022) View paper
    - Occlusion Recovery in Particle Objects (1 papers)
    - [15] Occlusion segmentation: restore and segment invisible areas for particle objects (Jinshi Liu, 2024) View paper
    - Vision-Language Model Analysis (1 papers)
    - [47] Mutual exclusivity bias and spatial reasoning in Vision-Language Models (Thai, 2025) View paper

Narrative

Core task: 3D part amodal segmentation with complete geometry inference. This field addresses the challenge of segmenting object parts in 3D while inferring their complete geometry even when portions are occluded or missing. The taxonomy reveals a rich landscape organized around several complementary perspectives. Amodal Completion and Shape Reconstruction Methods focus on recovering full object shapes from partial observations, often leveraging diffusion models or neural implicit representations to hallucinate occluded regions. Part-Level Segmentation and Decomposition emphasizes breaking objects into semantic components, with works like PartSAM[4] and Omnipart[6] exploring how to identify and delineate individual parts. Meanwhile, 2D Amodal Segmentation and Image-Based Methods tackle occlusion reasoning in image space, and Instance Segmentation and Tracking with Occlusion Handling extends these ideas to temporal scenarios. Geometry-Guided Segmentation and Reconstruction exploits geometric priors to constrain and improve predictions, while Domain-Specific Applications demonstrate how these techniques adapt to specialized contexts such as agriculture or medical imaging.

Recent activity has concentrated on diffusion-based approaches for amodal reconstruction, where generative models learn to complete missing geometry in a probabilistically principled manner. HoloPart[0] sits within this Diffusion-Based Amodal Reconstruction cluster, alongside works like pix2gestalt[1], Amodal3r[8], and AmodalGen3D[10], all of which harness diffusion priors to infer occluded part structures. Compared to Sequential Amodal[27], which may adopt iterative refinement strategies, HoloPart[0] emphasizes end-to-end generation of complete part geometries. The interplay between part-level decomposition methods such as Stable Part Diffusion[5] and holistic shape completion remains an open question: whether to first segment then complete, or to jointly infer parts and their full

extents. Domain-specific challenges, illustrated by works on occluded fruits or tree branches, highlight the need for robust occlusion handling across varied real-world scenarios, situating HoloPart[0] as part of a broader effort to unify geometric reasoning with modern generative modeling.

Related Works in Same Category

The following **4 sibling papers** share the same taxonomy leaf node with the original paper:

1. pix2gestalt: Amodal segmentation by synthesizing wholes

Authors: Ege Ozguroglu, Ruoshi Liu, DÃdac SurÃs, Dian Chen, D'idac Sur'is, et al. (9 authors total) | **Year/Venue**: 2024 | **URL**: View paper

Abstract

We introduce pix2gestalt, a framework for zero-shot amodal segmentation, which learns to estimate the shape and appearance of whole objects that are only partially visible behind occlusions. By capitalizing on large-scale diffusion models and transferring their representations to this task, we learn a conditional diffusion model for reconstructing whole objects in challenging zero-shot cases, including examples that break natural and physical priors, such as art. As training data, we use a synth...

Relationship Analysis

Both papers belong to the Diffusion-Based Amodal Reconstruction category, using diffusion models to generate complete shapes from partial inputs. While pix2gestalt focuses on 2D amodal segmentation by synthesizing whole objects from occluded 2D images using large-scale diffusion models, HoloPart addresses 3D part amodal segmentation by completing individual 3D part geometries within a larger shape using specialized local and context-aware attention mechanisms. The key difference is the dimensionality (2D vs 3D) and granularity (whole objects vs individual parts within objects).

---

2. Amodal3r: Amodal 3d reconstruction from occluded 2d images

Authors: Wu, Tianhao, Zheng, Chuanxia, Guan Frank, et al. (9 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Most image-based 3D object reconstructors assume that objects are fully visible, ignoring occlusions that commonly occur in real-world scenarios. In this paper, we introduce Amodal3R, a conditional 3D generative model designed to reconstruct 3D objects from partial observations. We start from a "foundation" 3D generative model and extend it to recover plausible 3D geometry and appearance from occluded objects. We introduce a mask-weighted multi-head cross-attention mechanism followed by an occlu...

Relationship Analysis

Both papers belong to the diffusion-based amodal reconstruction category, using diffusion models to generate complete 3D shapes from partial or occluded inputs. While HoloPart focuses on completing individual semantic parts within a 3D shape (part-level amodal segmentation) using a two-stage approach with local and context-aware attention mechanisms, Amodal3R addresses object-level amodal 3D reconstruction from occluded 2D images, introducing mask-weighted cross-attention and occlusion-aware layers to handle multi-view occlusions. The key distinction is that HoloPart operates on already-segmented 3D parts to complete their geometry, whereas Amodal3R reconstructs entire 3D objects from 2D occluded images without requiring prior segmentation.

---

3. AmodalGen3D: Generative Amodal 3D Object Reconstruction from Sparse Unposed Views

Authors: Junwei Zhou, Yu-Wing Tai | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Reconstructing 3D objects from a few unposed and partially occluded views is a common yet challenging problem in real-world scenarios, where many object surfaces are never directly observed. Traditional multi-view or inpainting-based approaches struggle under such conditions, often yielding incomplete or geometrically inconsistent reconstructions. We introduce AmodalGen3D, a generative framework for amodal 3D object reconstruction that infers complete, occlusion-free geometry and appearance from...

Relationship Analysis

Both papers belong to the Diffusion-Based Amodal Reconstruction category, using diffusion models to generate complete 3D geometry from partial or occluded inputs. While HoloPart focuses on completing individual 3D part segments within a whole shape using a two-stage approach with local and context-aware attention mechanisms, AmodalGen3D addresses the reconstruction of complete 3D objects from sparse, unposed, and occluded multi-view images by integrating 2D amodal completion priors with multi-view stereo geometry conditioning. The key distinction is that HoloPart operates on part-level completion within segmented shapes, whereas AmodalGen3D performs object-level reconstruction from occluded sparse views.

---

4. Sequential Amodal Segmentation via Cumulative Occlusion Learning

Authors: Ao, Jiayang, Jiayang Ao, Ke, Qiuhong, et al. (9 authors total) | **Year/Venue**: 2024 | **URL**: View paper

Abstract

To fully understand the 3D context of a single image, a visual system must be able to segment both the visible and occluded regions of objects, while discerning their occlusion order. Ideally, the system should be able to handle any object and not be restricted to segmenting a limited set of object classes, especially in robotic applications. Addressing this need, we introduce a diffusion model with cumulative occlusion learning designed for sequential amodal segmentation of objects with uncerta...

Relationship Analysis

Both papers belong to the Diffusion-Based Amodal Reconstruction category, using diffusion models to generate complete object shapes from occluded inputs. The original paper (HoloPart) focuses on 3D part amodal segmentation, completing individual 3D part geometries within a whole shape using a dual attention mechanism and shape context awareness. The candidate paper addresses 2D sequential amodal segmentation with cumulative occlusion learning, predicting multiple layers of 2D masks in order, which is fundamentally different from HoloPart's 3D part completion task despite both leveraging diffusion models for amodal reconstruction.

Contributions Analysis

Overall novelty summary. ```json { "paragraphs": ["The paper introduces 3D part amodal segmentation—decomposing 3D shapes into complete, semantically meaningful parts even under occlusion—and proposes HoloPart, a diffusion-based model for part completion. Within the taxonomy, this work resides in the 'Diffusion-Based Amodal Reconstruction' leaf, which contains five papers total. This leaf sits under 'Amodal Completion and Shape Reconstruction Methods,' indicating a moderately active research direction focused on generative approaches to occlusion reasoning. The taxonomy reveals that diffusion-based methods represent one of several competing paradigms for amodal reconstruction, alongside VAE/GAN-based approaches and multi-view techniques.",

"The taxonomy structure shows neighboring leaves addressing related but distinct challenges: 'Other Generative Amodal Approaches' explores non-diffusion generative models, while 'Single-View Amodal Reconstruction' focuses on RGB-based inference without the part-level decomposition emphasis. The 'Part Segmentation with Amodal Reasoning' leaf under a separate branch addresses similar occlusion

handling but from a segmentation-first perspective rather than generative completion. The scope notes clarify that this work differs from 'Part Segmentation without Amodal Reasoning' by explicitly modeling occluded geometry, and from 'Part-Aware 3D Generation' by operating on existing shapes rather than generating from scratch.",

"Among 30 candidate papers examined, the contribution-level analysis reveals mixed novelty signals. The task formulation and benchmarks (Contribution 1) show no clear refutation across 10 candidates, suggesting this specific problem framing may be relatively underexplored. The HoloPart diffusion architecture (Contribution 2) similarly shows no refutation among 10 examined papers, indicating potential architectural novelty within the diffusion-based reconstruction space. However, the two-stage approach combining segmentation and completion (Contribution 3) found one refutable candidate among 10 examined, suggesting some prior work explores similar decomposition-then-completion pipelines, though the scale of overlap remains limited given the search scope.",

"Based on this limited analysis of 30 semantically similar papers, the work appears to occupy a moderately explored niche within diffusion-based 3D reconstruction. The task-level novelty seems stronger than the methodological approach, with the two-stage pipeline showing some precedent. The taxonomy context suggests this sits at the intersection of two active areas—diffusion-based generation and part-level reasoning—where the specific combination may offer incremental contributions. A more exhaustive search beyond the top-30 semantic matches woul

---

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

Contribution 1: 3D part amodal segmentation task and benchmarks

Description: The authors define a new task called 3D part amodal segmentation that decomposes 3D shapes into complete semantic parts rather than visible surface patches. They establish two evaluation benchmarks using the ABO and PartObjaverse-Tiny datasets to enable future research in this area.

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. Perceptual organization and recognition of indoor scenes from RGB-D images

URL: View paper

Brief Assessment

Indoor Scene Recognition[52] focuses on 2D amodal segmentation and completion of surfaces in RGB-D indoor scenes, not 3D part-level decomposition of individual objects into complete semantic parts as proposed in the original paper.

---

2. Learning Vector Quantized Shape Code for Amodal Blastomere Instance Segmentation

URL: View paper

Brief Assessment

Vector Quantized Blastomere[26] focuses on 2D amodal segmentation of blastomere cells in microscopy images, not 3D shape decomposition. The candidate addresses biomedical cell segmentation with translucent overlapping objects, while the original paper introduces 3D part amodal segmentation for decomposing 3D shapes into complete semantic parts.

---

3. Application of amodal segmentation for shape reconstruction and occlusion recovery in occluded tomatoes

URL: View paper

Brief Assessment

Occluded Tomatoes[11] focuses on 2D amodal segmentation for agricultural tomato images, not 3D shape decomposition into semantic parts. The domains and technical approaches are fundamentally different.

---

4. Amodal3r: Amodal 3d reconstruction from occluded 2d images

URL: View paper

Brief Assessment

Amodal3r[8] focuses on reconstructing complete 3D objects from occluded 2D images, not on decomposing 3D shapes into semantic parts. The tasks are fundamentally different: Amodal3r[8] performs image-to-3D reconstruction with occlusion handling, while the original paper addresses 3D shape decomposition into complete semantic parts.

---

5. Fig-nerf: Figure-ground neural radiance fields for 3d object category modelling

URL: View paper

Brief Assessment

Fig-nerf[51] focuses on 2D amodal segmentation for figure-ground separation in neural radiance fields, not 3D part-level semantic decomposition of shapes into complete parts.

---

6. Single Point, Full Mask: Velocity-Guided Level Set Evolution for End-to-End Amodal Segmentation

URL: View paper

Brief Assessment

Velocity-Guided Level Set[44] focuses on 2D amodal segmentation in images (recovering complete object shapes from occluded regions in photos), not 3D shape decomposition into semantic parts. The domains and technical approaches are fundamentally different.

---

7. 3D-Aware Instance Segmentation and Tracking in Egocentric Videos

URL: View paper

Brief Assessment

Egocentric Instance Tracking[3] focuses on instance segmentation and tracking in egocentric videos with amodal video object segmentation as a downstream application, not on decomposing 3D shapes into complete semantic parts or establishing benchmarks for 3D part amodal segmentation.

---

8. DreamArt: Generating Interactable Articulated Objects from a Single Image

URL: View paper

Brief Assessment

DreamArt[12] focuses on generating articulated objects from single images with part decomposition and articulation modeling, not on establishing the 3D part amodal segmentation task itself or creating evaluation benchmarks for this specific task.

---

9. Application of amodal segmentation for shape

URL: View paper

Brief Assessment

Amodal Shape Application[23] focuses on applying amodal segmentation for shape analysis using 3D tree models, not on decomposing general 3D shapes into complete semantic parts or establishing benchmarks for this task.

---

10. Perceiving the Invisible: Proposal-Free Amodal Panoptic Segmentation

URL: View paper

Brief Assessment

Proposal-Free Amodal Panoptic[34] focuses on 2D amodal panoptic segmentation for traffic scenes, predicting semantic labels and instance shapes in images. This is fundamentally different from the original paper's 3D part amodal segmentation task, which decomposes 3D shapes into complete semantic parts.

Contribution 2: HoloPart diffusion-based model for 3D part shape completion

Description: The authors introduce HoloPart, a diffusion-based model designed specifically for completing 3D part shapes. The model incorporates local attention to capture fine-grained part geometry and context-aware attention to maintain global shape consistency, while leveraging pretrained 3D generative priors to handle limited training data.

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. Neusdfusion: A spatial-aware generative model for 3d shape completion, reconstruction, and generation

URL: View paper

Brief Assessment

Neusdfusion[60] focuses on whole 3D shape generation, completion, and reconstruction using a spatial-aware generative model with tri-plane representations. It does not address the specific problem of 3D part shape completion with attention mechanisms for individual semantic parts within a larger shape, which is the core contribution of HoloPart.

---

2. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation

URL: View paper

Brief Assessment

Sdfusion[53] focuses on whole-object shape completion and multi-modal 3D generation, not on completing individual semantic parts within a larger shape while maintaining global consistency. The candidate addresses different technical challenges (multi-modal conditioning for complete objects) compared to the original's focus on part-level completion with local and context-aware attention mechanisms.

---

3. Locally attentional sdf diffusion for controllable 3d shape generation

URL: View paper

Brief Assessment

Locally Attentional SDF[56] focuses on sketch-conditioned 3D shape generation from 2D images, not on completing 3D part shapes from incomplete segments. The tasks and technical approaches are fundamentally different.

---

4. Topology-aware latent diffusion for 3d shape generation

URL: View paper

Brief Assessment

Topology-aware Latent Diffusion[58] focuses on generating complete 3D shapes with topological constraints using persistent homology, not on completing individual 3D part shapes from segmentation masks. The candidate addresses whole-shape generation with topology awareness, while the original addresses part-level completion within segmented objects.

---

5. 3DShape2VecSet: A 3D Shape Representation for Neural Fields and Generative Diffusion Models

URL: View paper

Brief Assessment

3DShape2VecSet[61] focuses on a general shape representation for neural fields and generative diffusion models across various 3D generation tasks, not specifically on 3D part shape completion with attention mechanisms for handling incomplete part segments.

---

6. MVDiffusion++: A Dense High-resolution Multi-view Diffusion Model for Single or Sparse-view 3D Object Reconstruction

URL: View paper

Brief Assessment

MVDiffusion++[59] focuses on multi-view image synthesis for 3D object reconstruction from images, not on 3D part shape completion from segmented meshes. The technical approaches and problem domains are fundamentally different.

---

7. SGCDiff: Sketch-Guided Cross-modal Diffusion Model for 3D shape completion

URL: View paper

Brief Assessment

SGCDiff[55] focuses on sketch-guided cross-modal diffusion for 3D shape completion, not specifically on 3D part shape completion with local and context-aware attention mechanisms as proposed in HoloPart.

---

8. Wonder3D: Single Image to 3D Using Cross-Domain Diffusion

URL: View paper

Brief Assessment

Wonder3D[57] focuses on single-image to 3D reconstruction using cross-domain diffusion for generating multi-view normal maps and color images, not on 3D part shape completion with attention mechanisms for handling incomplete part segments.

---

9. Dora: Sampling and Benchmarking for 3D Shape Variational Auto-Encoders

URL: View paper

Brief Assessment

Dora[62] focuses on improving VAE reconstruction quality for general 3D shapes through sampling strategies and attention mechanisms, not on part-specific shape completion or amodal segmentation tasks that HoloPart addresses.

---

10. Diffusion models for 3D generation: A survey

URL: View paper

Brief Assessment

Diffusion 3D Survey[54] is a survey paper covering diffusion models for 3D generation broadly. The retrieved fragments mention diffusion-based 3D shape completion tasks in general, but provide no specific evidence of prior work on diffusion models designed specifically for 3D part shape completion with local and context-aware attention mechanisms as proposed in HoloPart.

Contribution 3: Two-stage approach combining segmentation and part completion

Description: The authors propose a practical two-stage pipeline where existing 3D part segmentation methods provide initial incomplete surface patches, followed by their HoloPart model that completes these segments into full 3D parts with proper geometry and semantic consistency.

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. Salad: Part-level latent diffusion for 3d shape generation and manipulation

URL: View paper

Brief Assessment

Salad[67] focuses on part-level latent diffusion for shape generation and manipulation, not on combining existing segmentation methods with completion models. The original paper's two-stage pipeline leverages existing segmentation followed by their novel HoloPart completion model, while Salad[67] presents a cascaded diffusion framework operating on part-level implicit representations for generation tasks.

---

2. Single image 3D object reconstruction based on deep learning: A review

URL: View paper

Brief Assessment

Single Image Review[63] mentions a two-stage decomposition for 2.5D to 3D shape reconstruction, but focuses on single-image reconstruction rather than the original paper's specific pipeline of leveraging existing 3D part segmentation followed by HoloPart's diffusion-based completion model for 3D part amodal segmentation.

---

3. Reconstructing dual-phase nanometer scale grains within a pearlitic steel tip in 3D through 4D-scanning precession electron diffraction tomography and automated â□¦

URL: View paper

Brief Assessment

Pearlitic Steel Grains[68] focuses on reconstructing nanometer-scale grains in metallurgical materials using electron diffraction tomography, not on 3D shape segmentation and completion pipelines for digital objects.

---

4. TreePartNet: neural decomposition of point clouds for 3D tree reconstruction

URL: View paper

Brief Assessment

TreePartNet[65] focuses on tree reconstruction from point clouds using cylindrical decomposition, not general 3D part segmentation and completion. The two-stage approach in TreePartNet is specific to detecting branching structures and creating cylindrical representations for trees, which is fundamentally different from the general 3D part amodal segmentation task.

---

5. Axis-aligned height-field block decomposition of 3D shapes

URL: View paper

Brief Assessment

Height-field Block Decomposition[70] focuses on decomposing 3D shapes into axis-aligned height-field blocks for reconstruction purposes, not on completing occluded semantic parts from segmentation masks as in the original paper's two-stage pipeline.

---

6. RANSAC-based multi primitive building reconstruction from 3D point clouds

URL: View paper

Brief Assessment

RANSAC Building Reconstruction[64] focuses on primitive-based building reconstruction from point clouds using RANSAC segmentation, not on learning-based 3D part shape completion with diffusion models for general objects.

---

7. Cylindrical Shape Decomposition for 3D Segmentation of Tubular Objects

URL: View paper

Brief Assessment

Cylindrical Shape Decomposition[71] focuses on decomposing tubular structures (axons, blood vessels) using curve skeletons and translational sweeps, not on general 3D part segmentation with shape completion for diverse object categories.

---

8. PartGen: Part-level 3D Generation and Reconstruction with Multi-View Diffusion Models

URL: View paper

Prior Art Analysis

PartGen[7] demonstrates that a two-stage pipeline combining part segmentation followed by part completion was already proposed prior to the original paper. Both papers describe extracting initial part segmentations (surface patches) in the first stage, then completing these segments into full 3D parts in the second stage. PartGen[7] explicitly states it uses 'a multi-view diffusion model to extract plausible and view-consistent part segmentations' followed by 'a second multi-view diffusion model then processes each part individually, filling in occlusions and generating completed views', which directly parallels the original paper's two-stage approach of 'part segmentation' followed by '3d part shape completion given segmentation masks'.

Evidence

Evidence 1 - **Rationale**: Both papers describe a two-stage pipeline where the first stage performs part segmentation to obtain initial surface patches, and the second stage completes these segments into full 3D parts. PartGen[7] uses multi-view diffusion models for both stages, demonstrating this approach existed prior to the original paper's submission. - **Original**: recognizing the inherent difficulty of end-to-end learning for this task, we propose a practical and effective two-stage approach. the first stage,part segmentation, has been widely studied, and we leverage an existing state-of-the-art method yang et al. (2024) to obtain initial, incomplete part seg... - **Candidate**: we introduce partgen, a novel approach for generating, from text, images, or unstructured 3d objects, 3d objects composed of meaningful

parts. our method leverages a multi-view diffusion model to extract plausible and view-consistent part segmentations from multiple views of a 3d object, dividing it...

Evidence 2 - **Rationale**: This evidence pair shows that PartGen[7] also addresses the challenge of completing occluded parts while maintaining global consistency, which is a key aspect of the original paper's two-stage approach. The candidate explicitly mentions 'compensating for missing information caused by occlusions' in the completion stage. - **Original**: to address these challenges, we propose a two-stage approach: 1.part segmentation:we first obtain an initial part segmentation of the input shapem. this provides us with a set of surface patches, each corresponding to a (potentially occluded) under review as a conference paper at iclr 2026 cross att... - **Candidate**: the completion process ensures that the reconstructed parts integrate cohesively by considering the context of the entire object, compensating for missing information caused by occlusions and, in extreme cases, hallucinating entirely invisible parts based on contextual cues.

---

9. Registration between DCT and EBSD datasets for multiphase microstructures

URL: View paper

Brief Assessment

DCT EBSD Registration[69] focuses on materials characterization and registration between diffraction-contrast tomography and electron backscatter diffraction datasets for multiphase microstructures, not on 3D part segmentation and shape completion pipelines for computer graphics applications.

---

10. Dual-phase airway segmentation: Enhancing distal bronchial identification with anatomical prior guidance

URL: View paper

Brief Assessment

Dual-phase Airway[66] focuses on airway segmentation in medical imaging with dual-decoding architecture, not on 3D part segmentation and shape completion for general 3D objects.

Appendix: Text Similarity Detection

Textual similarity detection checked 33 papers and found 2 similarity segment(s) across 2 paper(s).

The following **2 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. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation

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.

---

2. Wonder3D: Single Image to 3D Using Cross-Domain Diffusion

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] HoloPart: Generative 3D Part Amodal Segmentation View paper
- [1] pix2gestalt: Amodal segmentation by synthesizing wholes View paper
- [2] Large-Scale 3D Shape Reconstruction and Segmentation from ShapeNet Core55 View paper
- [3] 3D-Aware Instance Segmentation and Tracking in Egocentric Videos View paper
- [4] PartSAM: A Scalable Promptable Part Segmentation Model Trained on Native 3D Data View paper
- [5] Stable Part Diffusion 4D: Multi-View RGB and Kinematic Parts Video Generation View paper
- [6] Omnipart: Part-aware 3d generation with semantic decoupling and structural cohesion View paper
- [7] PartGen: Part-level 3D Generation and Reconstruction with Multi-View Diffusion Models View paper
- [8] Amodal3r: Amodal 3d reconstruction from occluded 2d images View paper
- [9] Obscured tree branches segmentation and 3D reconstruction using deep learning and geometrical constraints View paper
- [10] AmodalGen3D: Generative Amodal 3D Object Reconstruction from Sparse Unposed Views View paper
- [11] Application of amodal segmentation for shape reconstruction and occlusion recovery in occluded tomatoes View paper
- [12] DreamArt: Generating Interactable Articulated Objects from a Single Image View paper
- [13] Viser: Video-specific surface embeddings for articulated 3d shape reconstruction View paper
- [14] 3D Face Reconstruction with the Geometric Guidance of Facial Part Segmentation View paper
- [15] Occlusion segmentation: restore and segment invisible areas for particle objects View paper
- [16] BANG: Dividing 3D Assets via Generative Exploded Dynamics View paper
- [17] Unsupervised 3D Object Segmentation of Point Clouds by Geometry Consistency View paper
- [18] Imagine: Image-Guided 3D Part Assembly with Structure Knowledge Graph View paper
- [19] Semantic Amodal Segmentation View paper
- [20] Variational amodal object completion View paper
- [21] Monogrnet: A geometric reasoning network for monocular 3d object localization View paper
- [22] Multi-vision-based Localization and Pose Estimation of Occluded Apple Fruits for Harvesting Robots View paper
- [23] Application of amodal segmentation for shape View paper
- [24] PLUG: Revisiting Amodal Segmentation with Foundation Model and Hierarchical Focus View paper
- [25] GEOMetrics: Exploiting Geometric Structure for Graph-Encoded Objects View paper
- [26] Learning Vector Quantized Shape Code for Amodal Blastomere Instance Segmentation View paper
- [27] Sequential Amodal Segmentation via Cumulative Occlusion Learning View paper
- [28] Development of a Precise Tree Structure from LiDAR Point Clouds View paper
- [29] GIN: Generative invariant shape prior for amodal instance segmentation View paper
- [30] ASIA: Adaptive 3D Segmentation using Few Image Annotations View paper
- [31] NeuralDiff: Segmenting 3D objects that move in egocentric videos View paper
- [32] Part2Point: A Part-Oriented Point Cloud Reconstruction Framework View paper

- [33] Self-supervised Amodal Video Object Segmentation View paper
- [34] Perceiving the Invisible: Proposal-Free Amodal Panoptic Segmentation View paper
- [35] Holistic Pose Graph: Modeling Geometric Structure among Objects in a Scene using Graph Inference for 3D Object Prediction View paper
- [36] KineDiff3D: Kinematic-Aware Diffusion for Category-Level Articulated Object Shape Reconstruction and Generation View paper
- [37] Poster 140: Patella Shape is Associated with ACL Injury and Changes in KOOS Following ACL Reconstruction View paper
- [38] BLADE: Box-Level Supervised Amodal Segmentation through Directed Expansion View paper
- [39] Amodal detection of 3d objects: Inferring 3d bounding boxes from 2d ones in rgb-depth images View paper
- [40] 3D Part Segmentation via Geometric Aggregation of 2D Visual Features View paper
- [41] Structure-Guided Segmentation for 3D Neuron Reconstruction. View paper
- [42] ImVoxelGNet: Image to voxels geometry-aware projection for multi-view RGB-based 3D object detection. View paper
- [43] 3D deep object recognition and semantic understanding for visually-guided robotic service View paper
- [44] Single Point, Full Mask: Velocity-Guided Level Set Evolution for End-to-End Amodal Segmentation View paper
- [45] ViTA-Seg: Vision Transformer for Amodal Segmentation in Robotics View paper
- [46] Generating Objects with Part-Articulation from a Single Image View paper
- [47] Mutual exclusivity bias and spatial reasoning in Vision-Language Models View paper
- [48] A2VIS: Amodal-Aware Approach to Video Instance Segmentation View paper
- [49] ENHANCED DETECTION OF 3D INDIVIDUAL TREES IN FORESTED AREAS USING AIRBORNE FULL-WAVEFORM LIDAR DATA BY COMBINING NORMALIZED CUTS WITH SPATIAL DENSITY CLUSTERING View paper
- [50] Live Reconstruction of Large-Scale Dynamic Outdoor Worlds View paper
- [51] Fig-nerf: Figure-ground neural radiance fields for 3d object category modelling View paper
- [52] Perceptual organization and recognition of indoor scenes from RGB-D images View paper
- [53] Sdfusion: Multimodal 3d shape completion, reconstruction, and generation View paper
- [54] Diffusion models for 3D generation: A survey View paper
- [55] SGCDiff: Sketch-Guided Cross-modal Diffusion Model for 3D shape completion View paper
- [56] Locally attentional sdf diffusion for controllable 3d shape generation View paper
- [57] Wonder3D: Single Image to 3D Using Cross-Domain Diffusion View paper
- [58] Topology-aware latent diffusion for 3d shape generation View paper
- [59] MVDiffusion++: A Dense High-resolution Multi-view Diffusion Model for Single or Sparse-view 3D Object Reconstruction View paper
- [60] Neusdfusion: A spatial-aware generative model for 3d shape completion, reconstruction, and generation View paper
- [61] 3DShape2VecSet: A 3D Shape Representation for Neural Fields and Generative Diffusion Models View paper
- [62] Dora: Sampling and Benchmarking for 3D Shape Variational Auto-Encoders View paper
- [63] Single image 3D object reconstruction based on deep learning: A review View paper
- [64] RANSAC-based multi primitive building reconstruction from 3D point clouds View paper
- [65] TreePartNet: neural decomposition of point clouds for 3D tree reconstruction View paper
- [66] Dual-phase airway segmentation: Enhancing distal bronchial identification with anatomical prior guidance View paper
- [67] Salad: Part-level latent diffusion for 3d shape generation and manipulation View paper
- [68] Reconstructing dual-phase nanometer scale grains within a pearlitic steel tip in 3D through 4D-scanning precession electron diffraction tomography and automated â¦ View paper
- [69] Registration between DCT and EBSD datasets for multiphase microstructures View paper
- [70] Axis-aligned height-field block decomposition of 3D shapes View paper
- [71] Cylindrical Shape Decomposition for 3D Segmentation of Tubular Objects View paper