Novelty Assessment Report

Paper: Unsupervised Representation Learning for 3D Mesh Parameterization with Semantic and Visibility Objectives
PDF URL: https://openreview.net/pdf?id=9LYsvna4Sk
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2025-12-30

Abstract

Recent 3D generative models produce high-quality textures for 3D mesh objects. However, they commonly rely on the heavy assumption that input 3D meshes are accompanied by manual mesh parameterization (UV mapping), a manual task that requires both technical precision and artistic judgment. Industry surveys show that this process often accounts for a significant share of asset creation, creating a major bottleneck for 3D content creators. Moreover, existing automatic methods often ignore two perceptually important criteria: (1) semantic awareness (UV charts should align semantically similar 3D parts across shapes) and (2) visibility awareness (cutting seams should lie in regions unlikely to be seen). To overcome these shortcomings and to automate the mesh parameterization process, we present an unsupervised differentiable framework that augments standard geometry-preserving UV learning with semantic- and visibility-aware objectives. For semantic-awareness, our pipeline (i) segments the mesh into semantic 3D parts, (ii) applies an unsupervised learned per-part UV-parameterization backbone, and (iii) aggregates per-part charts into a unified UV atlas. For visibility-awareness, we use ambient occlusion (AO) as an exposure proxy and back-propagate a soft differentiable AO-weighted seam objective to steer cutting seams toward occluded regions. By conducting qualitative and quantitative evaluations against state-of-the-art methods, we show that the proposed method produces UV atlases that better support texture generation and reduce perceptible seam artifacts compared to recent baselines. We will make our implementation code publicly available upon acceptance of the paper.

Core Task Landscape

This paper addresses: **Unsupervised 3D Mesh UV Parameterization with Semantic and Visibility Awareness**

A total of **2 papers** were analyzed and organized into a taxonomy with **3 categories**.

Taxonomy Overview

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

- **Semantic-Aware UV Parameterization with Part-Based Decomposition**
- **Semantic Mesh Reconstruction with Integrated UV and Texture Generation**

Complete Taxonomy Tree

- Unsupervised 3D Mesh UV Parameterization with Semantic and Visibility Awareness Survey Taxonomy
- Semantic-Aware UV Parameterization with Part-Based Decomposition
  - Unsupervised Part-Based UV Learning with Visibility Optimization ★ (1 papers)
  - [0] Unsupervised Representation Learning for 3D Mesh Parameterization with Semantic and Visibility Objectives (Anon et al., 2026) View paper
  - Aligned UV Embedding for Cross-Shape Texture Tasks (1 papers)
  - [1] AUV-Net: Learning Aligned UV Maps for Texture Transfer and Synthesis (Zhiqin Chen, 2022) View paper
- Semantic Mesh Reconstruction with Integrated UV and Texture Generation
  - Human-Specific Semantic Mesh Reconstruction with Textures (1 papers)
  - [2] Semantic Human Mesh Reconstruction with Textures (Xiaoyu Zhan, 2024) View paper

Narrative

Core task: unsupervised 3D mesh UV parameterization with semantic and visibility awareness. The field addresses the challenge of automatically unwrapping 3D surfaces onto 2D texture space in ways that respect both geometric structure and semantic meaning. The taxonomy reveals two main branches. The first, Semantic-Aware UV Parameterization with Part-Based Decomposition, focuses on methods that explicitly segment meshes into meaningful parts—such as limbs or facial regions—and then optimize UV layouts to minimize distortion while preserving part boundaries and visibility constraints. The second branch, Semantic Mesh Reconstruction with Integrated UV and Texture Generation, takes a more holistic view by jointly learning mesh geometry, UV coordinates, and texture maps, often leveraging neural rendering or generative models to ensure consistency across these representations. Together, these branches reflect a shift from purely geometric unwrapping toward semantically informed strategies that better support downstream tasks like texture painting and appearance editing.

Within the part-based decomposition branch, a handful of works explore unsupervised learning of UV atlases that account for occlusion and part semantics without manual annotation. Unsupervised Mesh Parameterization[0] exemplifies this direction by combining visibility optimization with part-aware clustering, aiming to produce clean seams and efficient texture packing. In contrast, AUV-Net[1] emphasizes learning-based distortion minimization but does not explicitly enforce semantic part structure, while Semantic Human Mesh[2] integrates semantic labels more directly into the reconstruction pipeline. The original paper sits squarely in the unsupervised part-based cluster, distinguishing itself by jointly addressing visibility and semantic coherence without requiring labeled training data. This positions it as a bridge between classical geometric parameterization and emerging semantic reconstruction approaches, offering a practical middle ground for applications that demand both interpretability and automation.

Related Works in Same Category

No sibling papers were found in the same taxonomy leaf. A taxonomy-subtopic-level comparison will be produced instead.

Taxonomy-Level Summary

Both subtopics address learning-based UV parameterization for 3D meshes but target different objectives. The original leaf focuses on unsupervised part-based decomposition with visibility-aware seam placement for individual shapes, while the sibling emphasizes learning aligned UV embeddings across multiple shapes to enable cross-shape texture applications like transfer and synthesis.

Similarities: - Both involve learning UV parameterization through neural/differentiable frameworks - Both aim to improve UV mapping quality beyond traditional geometric methods - Both can potentially leverage semantic information (parts vs. cross-shape correspondence)

Differences: - Original leaf operates on single shapes with part segmentation and visibility optimization; sibling works across shape collections with alignment constraints - Original leaf emphasizes seam placement and visibility awareness; sibling prioritizes consistent UV spaces for texture transfer - Original leaf is explicitly unsupervised; sibling's supervision approach is unspecified but likely requires shape correspondence - Original leaf targets general UV quality with semantic parts; sibling targets specific applications (texture transfer/synthesis)

Suggested Search Directions: - Hybrid methods combining part-based visibility optimization with cross-shape alignment - Unsupervised approaches for learning aligned UV spaces without correspondence supervision - Visibility-aware texture transfer methods that consider occlusion across shape families

Sibling Subtopics

- **Aligned UV Embedding for Cross-Shape Texture Tasks** (leaves: 1, papers: 1)
- Scope: Methods learning aligned UV spaces across multiple shapes to enable texture transfer and synthesis applications.
- Exclude: Methods focused on single-shape reconstruction or without cross-shape alignment belong to reconstruction categories.

Contributions Analysis

Overall novelty summary. The paper proposes an unsupervised differentiable framework for UV parameterization that jointly optimizes for semantic part alignment and visibility-aware seam placement. According to the taxonomy, it occupies the 'Unsupervised Part-Based UV Learning with Visibility Optimization' leaf, which currently contains only this work and no sibling papers. This suggests the paper targets a relatively sparse research direction within the broader field of semantic-aware UV parameterization, where most prior efforts either focus on supervised methods or omit explicit visibility optimization.

The taxonomy reveals two main branches: part-based decomposition methods and integrated reconstruction pipelines. The original paper sits in the former, neighboring the 'Aligned UV Embedding for Cross-Shape Texture Tasks' leaf (one paper) and the 'Human-Specific Semantic Mesh Reconstruction' leaf (one paper). While cross-shape alignment methods prioritize texture transfer across multiple instances, and human-specific pipelines embed UV generation within full reconstruction, this work focuses narrowly on single-shape parameterization with semantic and visibility constraints. The taxonomy's scope notes clarify that methods without explicit part segmentation or visibility-aware seam optimization belong elsewhere, positioning this paper at the intersection of two design choices rarely combined in prior work.

Among 28 candidates examined, the semantic-aware objective (partition-and-parameterize strategy) shows one refutable candidate out of 10 examined, indicating some prior exploration of part-based UV learning. The unsupervised differentiable framework and visibility-aware objective each examined 9 candidates with zero refutations, suggesting these contributions face less direct overlap in the limited search scope. The statistics imply that while semantic part decomposition has precedent, the combination of unsupervised learning, differentiable optimization, and ambient-occlusion-driven seam placement appears less well-covered among the top-30 semantic matches retrieved.

Given the sparse taxonomy structure and limited search scope, the work appears to occupy a niche intersection of semantic awareness and visibility optimization. The analysis covers top-K semantic neighbors and does not claim exhaustive coverage of all UV parameterization literature. The single refutation for the semantic objective suggests incremental refinement rather than wholesale novelty, though the joint framework may still offer practical value for automating a labor-intensive manual process.

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

Contribution 1: Unsupervised differentiable framework for semantic- and visibility-aware UV parameterization

Description: The authors propose a two-stage framework that extends geometry-preserving UV parameterization with two novel perceptual objectives: semantic awareness (aligning UV charts with meaningful 3D parts) and visibility awareness (placing seams in less-visible regions). The framework is fully differentiable and trained end-to-end without supervision.

This contribution was assessed against **9 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Highâ□Fidelity Texture Transfer Using Multiâ□Scale Depthâ□Aware Diffusion

URL: View paper

Brief Assessment

Multiscale Depth Diffusion[15] focuses on texture transfer using diffusion models with depth guidance and multi-view consistency, not on UV parameterization learning with semantic and visibility objectives.

---

2. Neural Jacobian Fields: Learning Intrinsic Mappings of Arbitrary Meshes

URL: View paper

Brief Assessment

Neural Jacobian Fields[18] focuses on learning intrinsic mappings through jacobian fields for general mesh deformation tasks. While it addresses UV parameterization as one application, it does not incorporate semantic-awareness or visibility-awareness objectives, which are the core novelty claims of the original paper.

---

3. Weakly supervised joint transfer and regression of textures for 3-D human reconstruction

URL: View paper

Brief Assessment

Weakly Supervised Texture[13] focuses on texture transfer and regression for 3D human reconstruction from single images, not on UV parameterization learning with semantic and visibility objectives. The candidate addresses texture generation for invisible body parts, while the original paper proposes a framework for learning UV mappings with perceptual objectives.

---

4. Neural jacobian fields

URL: View paper

Brief Assessment

Neural Jacobian Fields[19] focuses on learning piecewise linear mappings via neural networks for UV parameterization but does not address semantic awareness (aligning UV charts with meaningful 3D parts) or visibility awareness (placing seams in less-visible regions). The candidate's approach is optimization-based and operates in the intrinsic gradient domain, whereas the original paper explicitly introduces novel perceptual objectives for semantic and visibility awareness that are absent in this candidate.

---

5. Vehicle reconstruction and texture estimation using deep implicit semantic template mapping

URL: View paper

Brief Assessment

Vehicle Semantic Template[14] focuses on vehicle reconstruction using implicit semantic template mapping for 3D geometry and texture from single images. It does not address UV parameterization, seam placement, or texture atlas generation for meshes.

---

6. UV Mapping with Graph Learning

URL: View paper

Brief Assessment

UV Graph Learning[17] focuses on supervised learning for semantic UV mapping using graph neural networks, not an unsupervised differentiable framework. The candidate explicitly states 'supervised graph learning framework for semantic uv mapping' and uses labeled training data, whereas the original proposes unsupervised learning with semantic and visibility objectives.

---

7. Facerefiner: High-fidelity facial texture refinement with differentiable rendering-based style transfer

URL: View paper

Brief Assessment

FaceRefiner[12] focuses on facial texture refinement using style transfer and differentiable rendering for UV maps, not on learning UV parameterization itself with semantic or visibility objectives.

---

8. Model-based Self-supervision for Dense Face Alignment and 3D Reconstruction

URL: View paper

Brief Assessment

Model-based Face Alignment[16] focuses on 3D face reconstruction using differentiable least-squares model fitting and semantic segmentation supervision, not UV parameterization. The candidate addresses face alignment and 3D morphable models, while the original paper targets mesh UV mapping with semantic and visibility objectives.

---

9. AUV-Net: Learning Aligned UV Maps for Texture Transfer and Synthesis

URL: View paper

Brief Assessment

AUV-Net[1] focuses on learning aligned UV maps across multiple shapes for texture transfer/synthesis, not on visibility-aware seam placement or per-shape semantic partitioning as in the original work.

Contribution 2: Semantic-aware objective using partition-and-parameterize strategy

Description: The authors introduce a semantic-aware objective that segments meshes into semantic parts using shape diameter function, applies per-part UV parameterization, and aggregates results into a unified atlas. This ensures UV charts correspond to semantically coherent regions, simplifying texture editing and cross-shape correspondence.

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. Autocuts: simultaneous distortion and cut optimization for UV mapping

URL: View paper

Brief Assessment

AutoCuts[26] focuses on joint optimization of distortion and cut placement for UV mapping but does not address semantic-aware partitioning or semantic alignment of UV charts across shapes. The candidate's approach is fundamentally different from the original paper's partition-and-parameterize strategy using shape diameter function for semantic segmentation.

---

2. UV Parametrization via Topological Disk Segmentation of Surfaces

URL: View paper

Brief Assessment

Topological Disk Segmentation[28] focuses on ensuring topological disk equivalence through Voronoi decomposition and Euler characteristic verification, not on semantic-aware segmentation using shape diameter function for texture editing and cross-shape correspondence as in the original paper.

---

3. An Embeddable Implicit IUVD Representation for Part-Based 3D Human Surface Reconstruction

URL: View paper

Brief Assessment

Embeddable IUVD[23] focuses on 3D human surface reconstruction using SMPL UV maps for implicit occupancy functions, not on general mesh UV parameterization or semantic chart alignment for texture editing across arbitrary shapes.

---

4. PartUV: Part-Based UV Unwrapping of 3D Meshes

URL: View paper

Prior Art Analysis

PartUV[27] demonstrates that a partition-and-parameterize strategy for UV unwrapping was already established prior to the original paper's submission. PartUV[27] explicitly describes a 'part-based UV unwrapping pipeline' that 'combines high-level semantic part decomposition' with parameterization in a 'top-down recursive framework.' This directly parallels the original paper's claimed novel contribution of segmenting meshes into semantic parts and applying per-part UV parameterization. Both papers use semantic decomposition followed by per-part parameterization and aggregation, indicating that the core partition-and-parameterize approach was not first proposed by the original authors.

Evidence

Evidence 1 - **Rationale**: Both papers describe the same fundamental approach: decompose the mesh into semantic parts, then parameterize each part independently. PartUV[27] explicitly states this is a 'part-based uv unwrapping pipeline' with 'semantic part decomposition,' directly matching the original paper's partition-and-parameterize strategy. - **Original**: we design a partition-and-

parameterize strategy: (b) compute a per-vertex semantic partition of the mesh, (c) learn a geometry-preserving uv parameterization (sec. 3.2) independently for each semantic part to obtain per-part uv islands, and then aggregate and pack these islands into a unified uv atl... - **Candidate**: We introduce partuv, a part-based uv unwrapping pipeline that generates significantly fewer, part-aligned charts while maintaining low distortion. Built on top of a recent learning-based part decomposition method partfield, partuv combines high-level semantic part decomposition with novel geometric ...

Evidence 2 - **Rationale**: The three-stage process described in the original paper (partition, per-part parameterization, aggregation/ packing) is conceptually identical to PartUV[27]'s approach of semantic decomposition followed by parameterization and packing, demonstrating prior work using this strategy. - **Original**: our semantic-aware parameterization has three stages: (i) compute a per-vertex semantic partition of the input mesh using the shape diameter function (shdf) shapira et al. (2008); (ii) apply the base uv-parameterization backbone from sec. 3.2 independently to each semantic 3d part to obtain per-part... - **Candidate**: partuv combines high-level semantic part decomposition with novel geometric heuristics in a top-down recursive framework. it ensures each chart's distortion remains below a user-specified threshold while minimizing the total number of charts. the pipeline integrates and extends parameterization and ...

---

5. Semantic UV mapping to improve texture inpainting for indoor scenes

URL: View paper

Brief Assessment

Semantic UV Inpainting[21] focuses on improving texture inpainting for indoor scenes by leveraging semantic segmentation to separate structural elements (walls, floors) for better UV mapping. The original paper addresses general 3D mesh parameterization with semantic and visibility objectives across diverse shapes, not specifically indoor scene reconstruction or texture inpainting workflows.

---

6. Optcuts: Joint optimization of surface cuts and parameterization

URL: View paper

Brief Assessment

OptCuts[22] focuses on joint optimization of surface cuts and parameterization to minimize distortion, not on semantic-aware partitioning. The candidate's mention of 'partitioned regions' refers to computational parallelization, not semantic segmentation for UV chart alignment.

---

7. Disentangled clothed avatar generation with layered representation

URL: View paper

Brief Assessment

Disentangled Clothed Avatar[20] focuses on layered representation for clothed avatar generation with component disentanglement (body, hair, clothes), not on UV chart alignment or mesh parameterization for texture editing.

---

8. Layout-aware single-image document flattening

URL: View paper

Brief Assessment

Document Flattening[25] focuses on rectifying 2D document images with layout-aware segmentation for dewarping, not 3D mesh UV parameterization with semantic part alignment.

---

9. Development of a building information model-guided post-earthquake building inspection framework using 3D synthetic environments

URL: View paper

Brief Assessment

Post-Earthquake Inspection[24] focuses on building inspection using synthetic environments and mentions UV mapping only as a technical process for 3D mesh visualization. It does not address semantic-aware UV parameterization, mesh segmentation strategies, or the partition-and-parameterize approach for texture editing.

---

10. AUV-Net: Learning Aligned UV Maps for Texture Transfer and Synthesis

URL: View paper

Brief Assessment

AUV-Net[1] aligns textures across shapes using a texture alignment module with basis decomposition, not mesh segmentation via shape diameter function followed by per-part parameterization as proposed in the original paper.

Contribution 3: Visibility-aware objective using ambient occlusion for seam placement

Description: The authors introduce a visibility-aware objective that uses ambient occlusion as a differentiable proxy for visual exposure. By backpropagating an AO-weighted seam loss, the method steers cutting seams toward occluded surface regions, reducing perceptible texture discontinuities after rendering.

This contribution was assessed against **9 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Automatic techniques for texture mapping in virtual urban environments

URL: View paper

Brief Assessment

Automatic Urban Texturing[11] focuses on texture map generation for urban building facades using depth maps to remove occlusions, not on UV parameterization or seam placement optimization using ambient occlusion.

---

2. Seamless texturing of archaeological data

URL: View paper

Brief Assessment

Seamless Archaeological Texturing[9] focuses on texture mapping for archaeological meshes using depth-map-based occlusion detection and color leveling, not on learning-based UV parameterization with differentiable ambient occlusion objectives for seam placement.

---

3. RoomPainter: View-Integrated Diffusion for Consistent Indoor Scene Texturing

URL: View paper

Brief Assessment

RoomPainter[5] focuses on indoor scene texture synthesis using diffusion models with multi-view consistency, not UV parameterization or seam placement optimization using ambient occlusion.

---

4. Vector-based occlusion detection for automatic facade texture mapping

URL: View paper

Brief Assessment

Vector Occlusion Detection[6] addresses occlusion detection for facade texture mapping in 3D city models, not UV parameterization or seam placement optimization. The candidate focuses on detecting occluded areas in building facades for texture compensation, while the original paper uses ambient occlusion as a differentiable loss term to guide seam placement in UV unwrapping.

---

5. UTexGen: Highâ□□quality texture reconstruction for largeâ□□scale scenes using multiâ□□view images

URL: View paper

Brief Assessment

UTexGen[8] addresses texture reconstruction for large-scale terrain using tile-based processing with global depth maps for occlusion detection. This is fundamentally different from the original paper's UV parameterization framework that uses ambient occlusion to guide cutting-seam placement in mesh unwrapping.

---

6. Uv mapping

URL: View paper

Brief Assessment

UV Mapping[3] mentions seam placement and ambient occlusion maps in a workflow context, but provides no technical detail about using AO as a differentiable proxy for visibility-aware seam optimization. The candidate appears to be a general tutorial or overview rather than a research contribution proposing this specific method.

---

7. Generating photorealistic urban digital twins with Occlusion-Aware texture mapping

URL: View paper

Brief Assessment

Photorealistic Urban Twins[7] focuses on city-scale texture mapping for urban digital twins with occlusion-aware methods, while the original paper addresses general 3D mesh UV parameterization with ambient occlusion-weighted seam objectives. The candidate's context is insufficient to determine if it uses ambient occlusion as a differentiable proxy for seam placement in the same manner.

---

8. Creating realistic clothing and armor for 3D game characters

URL: View paper

Brief Assessment

Realistic Game Characters[4] focuses on creating clothing and armor assets for game characters through modeling, texturing, and rigging workflows. It does not address UV parameterization, seam placement optimization, or ambient occlusion-based visibility objectives.

---

9. Ambient Occlusion-Using Ray-tracing and Texture Blending

URL: View paper

Brief Assessment

Ambient Occlusion Raytracing[10] focuses on computing ambient occlusion values for shadow rendering and texture blending to eliminate visible seams at texture borders. It does not address UV parameterization, seam placement optimization, or differentiable learning frameworks for steering cutting seams toward occluded regions as proposed in the original paper.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Unsupervised Representation Learning for 3D Mesh Parameterization with Semantic and Visibility Objectives View paper
- [1] AUV-Net: Learning Aligned UV Maps for Texture Transfer and Synthesis View paper
- [2] Semantic Human Mesh Reconstruction with Textures View paper
- [3] Uv mapping View paper
- [4] Creating realistic clothing and armor for 3D game characters View paper
- [5] RoomPainter: View-Integrated Diffusion for Consistent Indoor Scene Texturing View paper
- [6] Vector-based occlusion detection for automatic facade texture mapping View paper
- [7] Generating photorealistic urban digital twins with Occlusion-Aware texture mapping View paper
- [8] UTexGen: Highâ□quality texture reconstruction for largeâ□scale scenes using multiâ□view images View paper
- [9] Seamless texturing of archaeological data View paper
- [10] Ambient Occlusion-Using Ray-tracing and Texture Blending View paper
- [11] Automatic techniques for texture mapping in virtual urban environments View paper
- [12] Facerefiner: High-fidelity facial texture refinement with differentiable rendering-based style transfer View paper
- [13] Weakly supervised joint transfer and regression of textures for 3-D human reconstruction View paper
- [14] Vehicle reconstruction and texture estimation using deep implicit semantic template mapping View paper
- [15] Highâ□Fidelity Texture Transfer Using Multiâ□Scale Depthâ□Aware Diffusion View paper
- [16] Model-based Self-supervision for Dense Face Alignment and 3D Reconstruction View paper
- [17] UV Mapping with Graph Learning View paper
- [18] Neural Jacobian Fields: Learning Intrinsic Mappings of Arbitrary Meshes View paper
- [19] Neural jacobian fields View paper
- [20] Disentangled clothed avatar generation with layered representation View paper
- [21] Semantic UV mapping to improve texture inpainting for indoor scenes View paper
- [22] Optcuts: Joint optimization of surface cuts and parameterization View paper
- [23] An Embeddable Implicit IUVD Representation for Part-Based 3D Human Surface Reconstruction View paper
- [24] Development of a building information model-guided post-earthquake building inspection framework using 3D synthetic environments View paper

- [25] Layout-aware single-image document flattening View paper
- [26] Autocuts: simultaneous distortion and cut optimization for UV mapping View paper
- [27] PartUV: Part-Based UV Unwrapping of 3D Meshes View paper
- [28] UV Parametrization via Topological Disk Segmentation of Surfaces View paper