

Novelty Assessment Report

Paper: Auto-Regressive Surface Cutting

PDF URL: <https://openreview.net/pdf?id=9HeKCY11zl>

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-30

Abstract

Surface cutting is a fundamental task in computer graphics, with applications in UV parameterization, texture mapping, and mesh decomposition. However, existing methods often produce technically valid but overly fragmented atlases that lack semantic coherence. We introduce SeamGPT, an auto-regressive model that generates cutting seams by mimicking professional workflows. Our key technical innovation lies in formulating surface cutting as a next token prediction task: sample point clouds on mesh vertices and edges, encode them as shape conditions, and employ a GPT-style transformer to sequentially predict seam segments with quantized 3D coordinates. Our approach achieves exceptional performance on UV unwrapping benchmarks containing both manifold and non-manifold meshes, including artist-created, and 3D-scanned models. In addition, it enhances existing 3D segmentation tools by providing clean boundaries for part decomposition.

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: **Auto-Regressive Surface Cutting for Mesh UV Parameterization**

A total of **3 papers** were analyzed and organized into a taxonomy with **4 categories**.

Taxonomy Overview

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

- **Learning-Based Sequential Seam Generation**
- **Joint Optimization of Cuts and Parameterization**
- **Part-Based and Semantic UV Unwrapping**
- **Progressive Mesh Parameterization**

Complete Taxonomy Tree

- Auto-Regressive Surface Cutting for Mesh UV Parameterization Survey Taxonomy
- Learning-Based Sequential Seam Generation
 - Transformer-Based Auto-Regressive Cutting ★ (1 papers)
 - [0] Auto-Regressive Surface Cutting (Anon et al., 2026) [View paper](#)
- Joint Optimization of Cuts and Parameterization
 - Unified Cut-Distortion Optimization (1 papers)
 - [1] Optcuts: Joint optimization of surface cuts and parameterization (M Li, 2018) [View paper](#)
- Part-Based and Semantic UV Unwrapping
 - Part-Aligned Chart Generation (1 papers)
 - [2] PartUV: Part-Based UV Unwrapping of 3D Meshes (Zhaoning Wang, 2025) [View paper](#)
- Progressive Mesh Parameterization
 - Multi-Resolution Texture Mapping (1 papers)
 - [3] Texture mapping progressive meshes (Pedro V. Sander, 2001) [View paper](#)

Narrative

Core task: auto-regressive surface cutting for mesh UV parameterization. The field of mesh UV parameterization has evolved along several complementary directions, each addressing different aspects of the challenge of flattening 3D surfaces onto 2D texture space. Learning-Based Sequential Seam Generation explores data-driven approaches that predict where to cut meshes in a step-by-step fashion, often leveraging neural architectures to learn optimal seam placement from examples. Joint Optimization of Cuts and Parameterization seeks to simultaneously determine both the cut locations and the resulting UV layout, treating these traditionally separate stages as a unified problem; Optcuts[1] exemplifies this approach by co-optimizing distortion and seam quality. Part-Based and Semantic UV Unwrapping focuses on decomposing meshes according to meaningful geometric or semantic boundaries, as seen in PartUV[2], which leverages part-aware strategies to produce more intuitive texture atlases. Progressive Mesh Parameterization, illustrated by Texture Mapping Progressive Meshes[3], addresses parameterization in the context of level-of-detail hierarchies, ensuring that UV coordinates remain consistent across mesh simplifications.

A central tension across these branches involves the trade-off between automation and control: fully automatic methods may struggle with artist-desired layouts, while semantic or part-based approaches require additional annotations or segmentation. Another active theme is the balance between local geometric quality (minimizing distortion) and global atlas efficiency (reducing seam length and chart count). Auto-Regressive Surface Cutting[0] sits within the Learning-Based Sequential Seam Generation branch, specifically employing transformer-based architectures to predict seam edges in an auto-regressive manner. This positions it as a data-driven alternative to joint optimization methods like Optcuts[1], emphasizing learned sequential decision-making over explicit energy minimization, and contrasting with part-based strategies such as PartUV[2] by focusing on direct seam prediction rather than semantic decomposition.

Related Works in Same Category

No sibling papers and no sibling subtopics were found under the same parent taxonomy node; the paper appears structurally isolated in the taxonomy.

Contributions Analysis

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

Contribution 1: SeamGPT: Auto-regressive model for surface cutting seam generation

Description: The authors propose SeamGPT, which formulates surface cutting as a next-token prediction task using a GPT-style transformer to sequentially predict seam segments with quantized 3D coordinates, mimicking artist workflows for generating semantically meaningful cutting seams.

This contribution was assessed against **6 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. SeamCrafter: Enhancing Mesh Seam Generation for Artist UV Unwrapping via Reinforcement Learning

URL: [View paper](#)

Brief Assessment

SeamCrafter[7] builds upon and extends SeamGPT's autoregressive approach rather than refuting it. The candidate explicitly acknowledges SeamGPT as prior work and proposes enhancements through dual-branch encoding and DPO fine-tuning, indicating it is a follow-up improvement rather than evidence of prior art that would refute the original's novelty claim.

2. Automatic generation of NC cutter path from massive data points

URL: [View paper](#)

Brief Assessment

NC Cutter Path Generation[5] focuses on generating NC (numerical control) cutter paths from massive data points for manufacturing applications, not on auto-regressive models for 3D surface cutting seams or mesh parameterization tasks.

3. A strategy for finish cutting path generation of compound surfaces

URL: [View paper](#)

Brief Assessment

Finish Cutting Path[8] focuses on CNC machining path generation for compound surfaces using cutter location points, not auto-regressive transformer models for generating cutting seams on 3D meshes for UV parameterization.

4. Intelligent visual media processing: When graphics meets vision

URL: [View paper](#)

Brief Assessment

Graphics Meets Vision[6] is a survey paper covering intelligent visual media processing across graphics and vision. It does not propose auto-regressive models for seam generation or address surface cutting as a next-token prediction task.

5. An Improved Texture Mapping Model Based on Mesh Parameterization in 3D Garments

URL: [View paper](#)

Brief Assessment

Garment Mesh Parameterization[9] focuses on garment-specific mesh partitioning using geometrical feature lines and shortest-path methods, not auto-regressive transformer-based seam generation. The technical approaches are fundamentally different.

6. SEEK-CAD: A Self-REFINED GENERATIVE MODEL-ING FOR 3D PARAMETRIC CAD USING LOCAL INFER

URL: [View paper](#)

Brief Assessment

SEEK-CAD[10] focuses on parametric CAD model generation using LLMs with a self-refinement mechanism, not on auto-regressive surface cutting seam generation. The candidate addresses a completely different problem domain (CAD design automation) with different technical approaches (LLM-based code generation with visual feedback) compared to the original's mesh cutting and UV unwrapping task.

Contribution 2: Auto-regressive formulation of cutting seams as line segment sequences

Description: The authors introduce a novel formulation that represents cutting seams as ordered sequences of 3D line segments with quantized coordinates, enabling auto-regressive generation through a GPT-inspired transformer architecture that captures sequential dependencies in the cutting process.

This contribution was assessed against **1 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Image Background Filtering, Damage Detection and Location Registration in 3d Models for Bridge Inspection Using Image and Point Cloud Fusion

URL: [View paper](#)

Brief Assessment

Bridge Inspection Fusion[4] focuses on bridge inspection using image and point cloud fusion for damage detection and 3D registration, not on surface cutting or auto-regressive sequence generation for mesh processing.

Contribution 3: Seam-enhanced 3D part segmentation approach

Description: The authors demonstrate that SeamGPT-generated seams can enhance existing 3D part segmentation methods by providing clean part boundaries, addressing the limitation of blurry boundaries in current segmentation tools through a patch-based part segmentation methodology.

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. CutS3D: Cutting Semantics in 3D for 2D Unsupervised Instance Segmentation

URL: [View paper](#)

Brief Assessment

CutS3D[18] focuses on 2D unsupervised instance segmentation using 3D information to separate instances in natural images, not on refining 3D part segmentation boundaries. The candidate addresses a fundamentally different problem domain (2D instance segmentation from images) compared to the original paper's focus on 3D mesh part segmentation with clean boundaries.

2. Unsupervised Representation Learning for 3D Mesh Parameterization with Semantic and Visibility Objectives

URL: [View paper](#)

Brief Assessment

Unsupervised Mesh Parameterization[14] focuses on UV parameterization with semantic and visibility objectives for texture mapping workflows, not on refining 3D part segmentation boundaries using cutting seams as a primary contribution.

3. Flatten Anything: Unsupervised Neural Surface Parameterization

URL: [View paper](#)

Brief Assessment

Flatten Anything[19] focuses on neural surface parameterization for UV unwrapping and texture mapping, not on refining 3D part segmentation boundaries. While both papers generate cutting seams, their purposes differ fundamentally—Flatten Anything[19] uses seams for surface flattening, whereas the original paper uses seams to enhance segmentation boundary quality.

4. Mesh segmentation with concavity-aware fields

URL: [View paper](#)

Brief Assessment

Concavity-Aware Segmentation[15] focuses on mesh segmentation using concavity-sensitive fields to locate cutting boundaries, not on enhancing existing segmentation methods with clean boundaries as SeamGPT does.

5. Weld-seam identification and model reconstruction of remanufacturing blade based on three-dimensional vision

URL: [View paper](#)

Brief Assessment

Weld-seam Reconstruction[11] focuses on weld-seam identification for remanufacturing blades using point cloud segmentation and B-spline surface cutting, which is a domain-specific industrial application rather than a general 3D part segmentation methodology.

6. KnifeCut: Refining Thin Part Segmentation with Cutting Lines

URL: [View paper](#)

Brief Assessment

KnifeCut[12] focuses on 2D image segmentation refinement for thin structures using line-based interaction, not 3D part segmentation with cutting seams. The domains and technical approaches are fundamentally different.

7. FlexPara: Flexible Neural Surface Parameterization

URL: [View paper](#)

Brief Assessment

FlexPara[17] focuses on neural surface parameterization for UV unwrapping and mesh flattening, not on enhancing 3D part segmentation boundaries. While both papers involve cutting seams, FlexPara's seams serve parameterization purposes rather than refining segmentation boundaries as claimed in the original contribution.

8. Aircraft skin gap and flush measurement based on seam region extraction from 3D point cloud

URL: [View paper](#)

Brief Assessment

Aircraft Skin Gap Measurement[13] focuses on detecting and segmenting seam areas in aircraft skin for gap/flush measurement applications, not on enhancing general 3D part segmentation methods or addressing blurry boundary problems in segmentation tools.

9. Intelligent visual media processing: When graphics meets vision

URL: [View paper](#)

Brief Assessment

Graphics Meets Vision[6] discusses general 3D part segmentation methods but does not propose using cutting seams to refine segmentation boundaries or address the blurry boundary problem through seam-based approaches.

10. Segmentation-Based Seam Cutting for High-Resolution 360-Degree Video Stitching

URL: [View paper](#)

Brief Assessment

Seam Cutting Video Stitching[16] focuses on 2D video stitching for 360-degree videos using segmentation to find seam boundaries, not on 3D part segmentation or mesh decomposition tasks.

Appendix: Text Similarity Detection

Textual similarity detection checked 16 papers and found 2 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. SeamCrafter: Enhancing Mesh Seam Generation for Artist UV Unwrapping via Reinforcement Learning

Detected in: Contribution: contribution_1

△ **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] Auto-Regressive Surface Cutting [View paper](#)
- [1] Optcuts: Joint optimization of surface cuts and parameterization [View paper](#)

- [2] PartUV: Part-Based UV Unwrapping of 3D Meshes [View paper](#)
- [3] Texture mapping progressive meshes [View paper](#)
- [4] Image Background Filtering, Damage Detection and Location Registration in 3d Models for Bridge Inspection Using Image and Point Cloud Fusion [View paper](#)
- [5] Automatic generation of NC cutter path from massive data points [View paper](#)
- [6] Intelligent visual media processing: When graphics meets vision [View paper](#)
- [7] SeamCrafter: Enhancing Mesh Seam Generation for Artist UV Unwrapping via Reinforcement Learning [View paper](#)
- [8] A strategy for finish cutting path generation of compound surfaces [View paper](#)
- [9] An Improved Texture Mapping Model Based on Mesh Parameterization in 3D Garments [View paper](#)
- [10] SEEK-CAD: A Self-REFINED GENERATIVE MODEL-ING FOR 3D PARAMETRIC CAD USING LOCAL INFER [View paper](#)
- [11] Weld-seam identification and model reconstruction of remanufacturing blade based on three-dimensional vision [View paper](#)
- [12] KnifeCut: Refining Thin Part Segmentation with Cutting Lines [View paper](#)
- [13] Aircraft skin gap and flush measurement based on seam region extraction from 3D point cloud [View paper](#)
- [14] Unsupervised Representation Learning for 3D Mesh Parameterization with Semantic and Visibility Objectives [View paper](#)
- [15] Mesh segmentation with concavity-aware fields [View paper](#)
- [16] Segmentation-Based Seam Cutting for High-Resolution 360-Degree Video Stitching [View paper](#)
- [17] FlexPara: Flexible Neural Surface Parameterization [View paper](#)
- [18] CutS3D: Cutting Semantics in 3D for 2D Unsupervised Instance Segmentation [View paper](#)
- [19] Flatten Anything: Unsupervised Neural Surface Parameterization [View paper](#)