SOVEREIGN: How does the inference latency of the proposed DDRNet23-slim variant compare to baseline segmentation models o

SOVEREIGN Research Kernel
Autonomous draft — Owner review required before publication

May 28, 2026

Abstract

Point clouds provide a flexible geometric representation suitable for countless applications in computer graphics; they also comprise the raw output of most 3D data acquisition devices. While hand-designed features on point clouds have long been proposed in graphics and vision, however, the recent overwhelming success of convolutional neural networks (CNNs) for image analysis suggests the value of adapting insight from CNN to the point cloud world. Point clouds inherently lack topological information, so designing a model to recover topology can enrich the representation power of point clouds.

1 Introduction

Analysis of: Dynamic Graph CNN for Learning on Point Clouds. Research goal: How does the inference latency of the proposed DDRNet23-slim variant compare to baseline segmentation models on the RUGD dataset when evaluated using mean intersection over union (mIoU) and frames per second (FPS) metrics?.

2 Methodology

Multi-query arXiv search (4 parallel queries, Relevance-sorted). TF-IDF cosine semantic verification (bigrams, threshold=0.15). NIM nv-embedqa-e5-v5 (dim=1024) for semantic indexing. Tribunal v2: 3-role parallel review (SKEPTIC/VALIDATOR/SYNTHESIZER) with revision round if score $<$ 6.5.

3 Results

9 papers retrieved. 0 claims extracted, 0 verified. Tribunal: 7.0/10 → REVISE (revision_round=1). Policy: ESCALATE_TO_OWNER.

4 Uncertainties

NIM free tier latency varies. TF-IDF verification is a weak signal. arXiv Relevance ranking is query-dependent. Tribunal consensus is LLM-based and prompt-sensitive.

References

- https://doi.org/10.1145/3326362
- https://doi.org/10.1186/s40537-021-00444-8
- https://doi.org/10.1109/jstars.2022.3219724