

SOVEREIGN: Can a mixture-of-experts (MoE) routing strategy with adaptive sparsity improve both throughput (FPS) and per-c

SOVEREIGN Research Kernel

Autonomous draft — Owner review required before publication

May 29, 2026

Abstract

This review comprehensively investigates the current state and emerging trends of autonomous vehicle terrain detection and segmentation. By systematically reviewing literature from various databases, this study outlines the evolution of detection and segmentation techniques from traditional computer vision methods to advanced machine learning and deep learning approaches. It identifies critical technological advancements, evaluates their performance, and discusses the challenges faced under various environmental conditions, data acquisition, and integration with vehicle systems. This study als

1 Introduction

Analysis of: Terrain detection and segmentation for autonomous vehicle navigation: A state-of-the-art systematic review. Research goal: Can a mixture-of-experts (MoE) routing strategy with adaptive sparsity improve both throughput (FPS) and per-class IoU for rare categories (vegetation, obstacles) compared to standard dense transformers on the KITTI-360 dataset?.

2 Methodology

Multi-query arXiv search (1 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

6 papers retrieved. 0 claims extracted, 0 verified. Tribunal: 3.7/10 \rightarrow REJECT (revision_round=0). 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/10.1145/3676536.3676681>
- <https://doi.org/10.1002/eom2.12194>
- <https://doi.org/10.1016/j.inffus.2024.102644>