Adversarially Trained Trajectory Prediction Models: Latency and Accuracy Trade-offs in Autonomous Driving

Assignee Research

May 30, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: How do adversarially trained trajectory prediction models compare in inference latency and accuracy trade-offs when evaluated on standard autonomous driving planning benchmarks. We introduce a motion forecasting (behavior prediction) method that meets the latency requirements for autonomous driving in dense urban environments without sacrificing accuracy. A whole-scene sparse input representation allows StopNet to scale to predicting trajectories for. 6 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 5.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: StopNet: Scalable Trajectory and Occupancy Prediction for Urban Autonomous Driving. Research question: How do adversarially trained trajectory prediction models compare in inference latency and accuracy trade-offs when evaluated on standard autonomous driving planning benchmarks?.

2 Methodology

Systematic literature search across multiple databases yielded 15 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 5.5/10.

3 Results

15 papers retrieved. 6 claims extracted; 1 independently verified. Quality review score: 5.5/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

5 Extracted Claims

Claim	Verified	Confidence
StopNet uses a sparse scene representation with 160$\times$160 pillars for optimal performance.	×	0.09
StopNet supports both trajectory and occupancy grid output formats.	×	0.10
Trajectory models output tens of potential trajectories per agent.	×	0.05
Occupancy grids require fixed compute to generate and consume regardless of the number of agents in the scene.	×	0.09
StopNet's whole-scene encoder supports predicting probabilistic occupancy grids.	✓	0.21
StopNet is the first method to unify trajectory sets and occupancy grids as the two archetypes of motion forecasting.	×	0.13

References

- http://arxiv.org/abs/2205.14230v2
- http://arxiv.org/abs/2206.00991v1
- http://arxiv.org/abs/2310.07794v2