Building Real-time Awareness of Out-of-distribution in Trajectory Prediction for Autonomous Vehicles

Tongfei Guo[1]
Department of Electrical & Computer Engineering,
Northeastern University, USA
guo.t@northeastern.edu

Taposh Banerjee [2]
Swanson School of Engineering,
University of Pittsburgh, USA
taposh.banerjee@pitt.edu

Lili Su[1]
Department of Electrical & Computer Engineering,
Northeastern University, USA
l.su@northeastern.edu

How do AVs stay safe in unpredictable conditions?

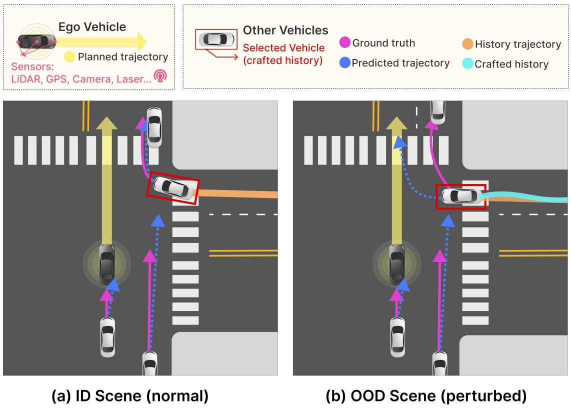

Figure 1: Performance comparison of a trajectory prediction model in ID and OOD scenes. In ID scenarios (left), the ego vehicle accurately predicts neighboring trajectories. In OOD scenarios (right), unexpected debris causes a slight deviation in the target vehicle's path, leading to incorrect trajectory prediction and potentially unsafe braking by the ego vehicle.

Motivation

Reliable autonomous driving requires trajectory prediction models to remain robust under real-world distribution shifts. Due to unavoidable sim-to-real gaps between training and inference, even well-trained models may produce unreliable predictions.

What challenges are involved?

- Deceptive OOD scenarios are difficult to detect by human intuition
- In OOD scene, non-trivial trajectory changes lead to significant prediction errors
- Traditional OOD detection focuses on single-point anomalies, overlooking sequential patterns
- Autonomous vehicles require real-time, sequential decision-making

Contribution

By formulating out-of-distribution detection as a quickest change-point detection problem, our approach enables timely identification of subtle and deceptive shifts in driving scenes. Our approach monitors only a scalar error variable, handles OOD occurrence at any inference step, and remains computationally efficient. We are the first to apply QCD methods for OOD detection in trajectory prediction across multiple real-world datasets.

Methods

Sequential analysis

- Quickest change-point detection (QCD)
  - CUSUM (selected)
    - knowledge level (ours)
      - Complete → CUSUM-Mix
      - Partial → CUSUM-Sinmix
      - Unknown → CUSUM-Single
- Time-series detection (Benchmarks)

How CUSUM works?

CUSUM monitor cumulative log-likelihood ratio Wt of AV trajectory prediction errors and triggers an alarm when Wt>b.

Figure 2: Statistical evolution of CUSUM detection given the prediction errors from both ID and OOD scenes.

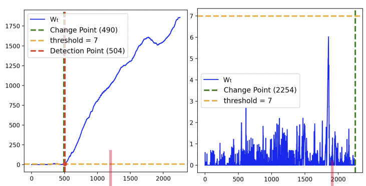

OOD scene: A perturbation at step 490 causes Wt to exceed the threshold b=7at step 507, yielding a 17-step detection delay.

ID scene: Prediction errors remain stable; Wt stays below b, indicating no false alarms.

Results

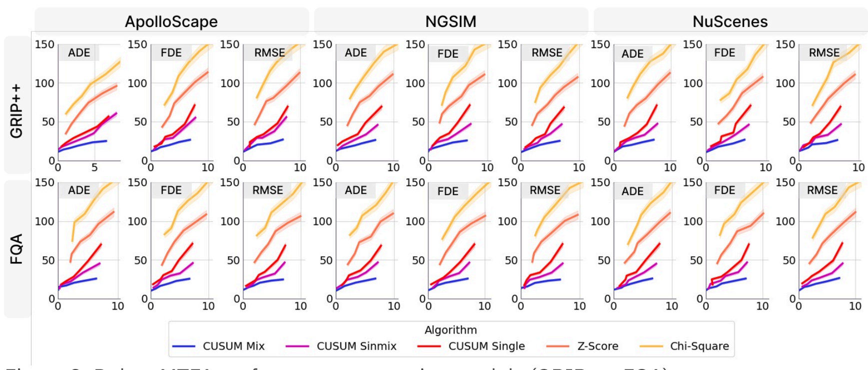

Figure 3: Delay–MTFA performance comparing models (GRIP++, FQA) across datasets (ApolloScape, NGSIM, NuScenes) and multiple metrics (ADE, FDE, RMSE).

Key Takeaways

- Lightweight QCD: Monitors a scalar prediction-error statistic.
- Real-world evaluation: Validated on ApolloScape, NGSIM, and NuScenes with GRIP++ and FQA.
- Effective detection:
  - Best performance: CUSUM Mix achieves lowest delay with minimal false alarms across all settings.
  - GMM benefit: GMM-based pre/post modeling improves robustness.
  - Delay–MTFA trade-off: Consistent advantage across metrics (Fig. 3).

Future Work

- Tiered alarm system: Introduce context-aware, multi-level alerts to prioritize critical warnings and reduce computational overhead.
- Adaptive learning: Leverage imitation learning to adapt detection behavior in OOD scenes.

AAAI-26 / IAAI-26 / EAAI-26

Northeastern University

Let's connect