DEPARTMENT OF ELECTRICAL & COMPUTER ENGINEERING
Cyber-Physical Systems Engineering Program

ENEB408A CAPSTONE DESIGN LAB 1 (Fall 2025)

RoadGuardian: Retrofit ADAS Using Computer Vision with Fuzzy Logic and Time-To-Collision Estimation

Jack Kouncar, Nour Hammad, Sebastian Yanez
Adviser: Dr. Nestor Tiglao

A. JAMES CLARK SCHOOL OF ENGINEERING

BACKGROUND

Problem

- Older vehicles often lack lane departure and forward-collision warnings.
- Modern ADAS is usually tied to expensive sensors, proprietary integration, or newer car models.
- Delayed reaction and limited awareness remain major crash-risk factors.

Road Guardian Approach

- Low-cost retrofit ADAS for vehicles without factory safety tech.
- Advisory only: warns without controlling steering, throttle, or brakes.
- Runs locally with one forward camera and OBD-II speed data.

Key contribution: A vehicle-agnostic, standalone device that combines computer vision, fuzzy logic, and TTC estimation to deliver practical driver alerts on affordable hardware.

MOTIVATION AND OBJECTIVES

Objectives

- Detect lane boundaries and lane-departure risk in real time.
- Detect lead vehicles using an optimized YOLO model.
- Estimate distance and Time-To-Collision (TTC).
- Keep setup plug-and-play, low-cost, and cloud-independent.

Hardware Prototype

Prototype Stack

- Raspberry Pi 5 with active cooling.
- 1080p forward-facing camera.
- Bluetooth ELM327 OBD-II adapter.
- Windshield/suction mounting hardware.
- Small display plus LED/buzzer alert outputs.
- Fully offline processing for low-latency feedback.

METHODOLOGY AND SYSTEM DESIGN

Prototype Data Flow

Perception and Decision Modules

Lane Detection + Fuzzy Logic

- OpenCV preprocessing with grayscale, edges, and adaptive HSV filtering.
- Fuzzy controller scores quality from lane count, point density, brightness, and noise.
- Gradual confidence changes replace brittle fixed thresholds.

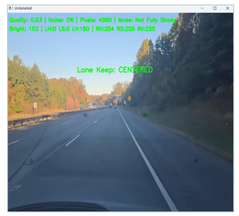

Vehicle Detection

- YOLOv11 detector trained on roughly 1,500 labeled vehicle images.
- ONNX deployment at 416 x 416 input resolution for Raspberry Pi CPU inference.
- Reuses detections between frames to reduce CPU load.

Distance + TTC Estimation

- Monocular distance is calibrated from detected vehicle bounding-box width.
- TTC combines distance with speed/closing behavior to set warning level.
- Safeguards suppress stale-speed and very-low-speed alerts.

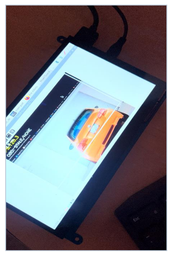

Driver Alerting and Deployment

- Visual overlays show lanes, vehicles, distance, TTC, and lane status.
- LED and buzzer outputs provide immediate driver feedback.
- Advisory design preserves driver control during testing.

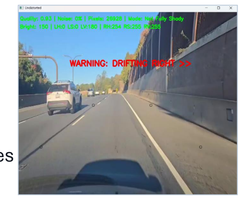

Implementation Details

Embedded Platform
- Linux-based Raspberry Pi 5
- Python + OpenCV runtime
- CPU-only ONNX inference
- Camera and OBD-II modules run in parallel

Adaptive Lane Quality
- Lane count: zero, one, two
- Point density: very low to high
- Brightness: dark, medium, bright
- Noise level: low, moderate, high

Alert Logic
- Lane centrality checked every frame
- TTC warnings triggered for closing risk
- False alerts reduced with speed and stale-data checks

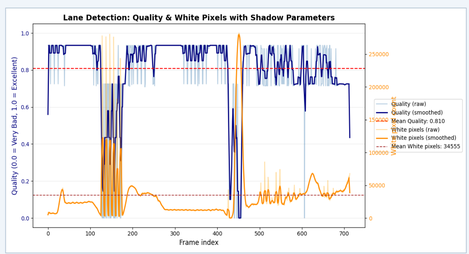

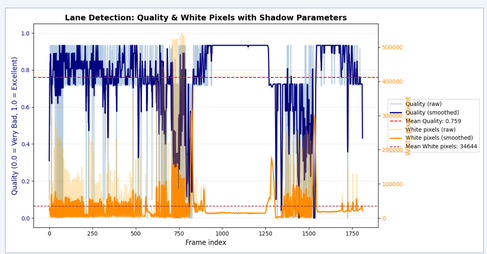

Fuzzy logic testing showed how detection quality changes as lane count, white-pixel count, and noise vary across driving conditions.

System goal: retrofit lane awareness and forward-collision warnings without expensive sensors, cloud services, or direct vehicle control.

RESULTS AND ANALYSIS

Measured Performance

97%	91.2%	96.2%	40-60
mAP@50 vehicle detection	precision	recall	FPS on Pi CPU

- YOLO vehicle detection ran in real time on Raspberry Pi 5 without GPU/external acceleration.
- Lane detection reached approximately 95% detection rate in daytime tests.
- Fuzzy quality assessment achieved about 85-90% accuracy and improved robustness to lighting/noise variation.

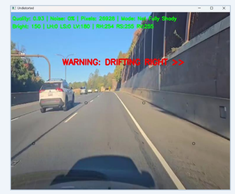

Example outputs: multi-vehicle detection with confidence overlays, and lane-departure warning when the vehicle drifts right.

Validation Summary

- **Tested with prerecorded road footage and real-world daytime driving.**
- **Alert triggering was stable and predictable during steady following scenarios.**
- **Fully local processing avoids network dependence and supports low-latency alerts.**

CONCLUSION AND FUTURE WORK

Road Guardian demonstrates that practical ADAS features can be retrofitted into older vehicles using affordable embedded hardware, a single monocular camera, OBD-II speed data, and adaptive computer vision.

Future Work
- Add Hailo AI acceleration for higher-resolution and larger YOLOv11 variants.
- Improve depth reliability with radar, ultrasonic, or low-light/IR sensors.
- Extend TTC with context-aware thresholds and lane-aware prioritization.
- Add mobile/web trip summaries and fleet-level safety dashboards.

References

[1] R. B. Naumann et al., AAA Foundation for Traffic Safety, 2023.
[2] S. Raviteja and R. Shanmughasundaram, ICICCS, 2018.
[3] G. S. P. K. Segu et al., real-time road lane and vehicle detection on YOLOv8, CICN, 2024.
[4] I.-C. Sang and W. R. Norris, robust lane detection for challenging weather, IEEE Access, 2024.
[5] G. Jocher and J. Qiu, Ultralytics YOLO11, 2024.