


Introduction







Autonomous driving aims to reclaim billions of person-hours and eliminate human errors responsible for over 94% of road fatalities. This thesis develops a complete autonomous driving pipeline on a one-tenth scale platform, reducing human labour in both *driving* and *system development* through AI-assisted tools and adaptive control. The system targets reliable **Level 4** autonomy via three pillars: **Perception**, **Planning**, and **Control**.




Perception

We developed a fully proprietary perception stack spanning annotation, FOV expansion, lane detection, and object detection.

LaneLab: Self-improving lane annotation engine using RL-based tuning and temporal backbone. Reduces annotation time by **85.6%** at 96.8% IoU. **DetLab**: Discover-name-learn object annotator—machine finds and boxes objects; human only names them. Reduces annotation time by **70%**. **SeamStitch**: RL-guided FOV expansion ($68^\circ \rightarrow 120^\circ$) via self-calibration and PPO-based extrinsic correction. 35.2 dB PSNR, 73 ms latency. **Aethris**: Lane detection on **Recochet** backbone with anisotropic bounce attention and geometric reconstruction for occluded lanes. **81.2% F1** at 82 FPS. **Point Blank**: Transformer-based detector with deformable cross-scale attention and occlusion-aware reconstruction. **96.2% mAP@0.5** at 38.9 FPS on edge hardware.





	Stop	Left Turn	Right Turn	Background
Predicted Stop	97.1	0.8	0.6	1.5
Predicted Left Turn	0.5	95.4	2.3	1.8
Predicted Right Turn	0.4	2.1	96.1	1.4
Predicted Background	2.0	1.7	1.0	95.3



Localization & Mapping


The vehicle must know *where it is* with centimeter-level accuracy using three complementary strategies: **SLAM**: Builds a map while tracking position in real time for navigation in unseen areas. **Dead Reckoning**: Wheel odometry and inertial measurements estimate displacement between absolute fixes. **Motion Capture**: Sub-millimetre ground-truth positioning at 120 Hz serves as both reference and validation benchmark.



Planning

Once localized, the vehicle decides *what to do* and *how to move* via a three-layer pipeline:

- Costmap Generation**: Transforms SLAM occupancy grid into a graded cost field with exponential decay around obstacles ($r_{\text{infl}} = 0.30 \text{ m}$).
- Global Planning (A*)**: Finds least-cost path evaluating $f(n) = g(n) + h(n)$, combining traversal distance with costmap penalty.
- Local Control (DWA)**: Samples candidate (v, ω) within dynamic window and selects the command maximizing heading alignment, obstacle clearance, and speed. Best route selection, lane changes with LiDAR and camera detected dynamic obstacles, and three-phase overtaking manoeuvres emerge naturally from A* replanning.





Control

Control converts the planned trajectory into physical vehicle motion at two levels.


Low-Level Control (Motor/Drivetrain)

Q-Learning-Optimised PID: Tabular Q-Learning agents discover optimal control laws offline. Each Q-table maps sensor states to PWM values via a single $O(1)$ read—deterministic sub- $3 \mu\text{s}$ latency on ARM Cortex-M7. Five regime-specific agents: **Zero-Jerk Launch** (93% jerk reduction), **Emergency Brake**, **ABS**, **Banked Surface**, and **Cruise Control**. All tables total only 7,792 bytes.



High-Level Control (Steering)


MPC: Predicts future states via bicycle model and solves receding-horizon optimisation minimising lateral/heading error subject to actuator limits. **Hybrid MPC-SAC**: SAC agent modulates MPC parameters (reference offset, weights, lane-change flag) rather than steering directly. Combines RL adaptability with model-based safety reducing collision rate, **98.7% pothole avoidance in controlled environment.**, 20 Hz real-time execution.



Safety & Validation

Human drivers has average one fatal crashes. Our multi-fidelity validation pipeline:

- 2D Simulation**: Rapid prototyping with controlled scenarios across hundreds of runs.
- 3D Simulation**: Full vehicle dynamics with domain randomisation over road, weather, and obstacle conditions.
- Motion Capture Lab**: Sub-millimetre tracking at 120 Hz for precise trajectory and controller comparison.
- Redundancy**: Deterministic state-machine agent switching, redundant sensor fusion, bounded sub- $3 \mu\text{s}$ latency, and **Minimal Risk Condition (MRC)** for safe stops.



Challenges & Future Directions

- Edge Cases**: Construction zones, emergency vehicles, and adversarial weather remain difficult.
- NeRF & 3D Gaussian Splatting**: Synthesise photorealistic views for rare scenario training.
- World Models**: Large-scale video prediction for generalised driving policies.
- Full-Scale Deployment**: Moving from 1/10th scale to real-world testing.
- V2X & Regulation**: Infrastructure adoption and evolving certification standards.

Conclusion

This work presents a fully autonomous 1/10th-scale vehicle unifying annotation-driven self-learning perception, motion-capture-validated localization, and Q-Learning/MPC-SAC control—achieving robust lane and object detection, 93% jerk reduction, zero collisions, and 98.7% pothole avoidance in real-time ultimately reducing human labor in collecting dataset and driving car and saving large chunk of time.