machines

MDPI

Article

A Novel Lateral Control System for Autonomous Vehicles: A Look-Down Strategy

Farzad Nadiri [1,2] and Ahmad B. Rad [1,*]

[1] Autonomous and Intelligent Systems Laboratory, School of Mechatronic Systems Engineering, Simon Fraser University, Surrey, BC V3T 0A3, Canada
[2] Quartech, 4211 Kingsway #710, Burnaby, BC V5H 1Z6, Canada
* Correspondence: arad@sfu.ca

Abstract: This paper introduces a robust yet straightforward lane detection and lateral control approach via the deployment of a dual camera based on the look-down strategy for autonomous vehicles. Unlike traditional single-camera systems that rely on the look-ahead methodology and a single front-facing preview, the proposed algorithm leverages two downward-facing cameras mounted beneath the vehicle's driver and the passenger side mirror, respectively. This configuration captures the road surface, enabling precise detection of the lateral boundaries, particularly during lane changes and in narrow lanes. A Proportional-Integral-Derivative (PID) controller is designed to maintain the vehicle's position in the center of the road. We compare this system's accuracy, lateral steadiness, and computational efficiency against (1) a conventional bird's-eye view lane detection method and (2) a popular deep learning-based lane detection framework. Experiments in the CARLA simulator under varying road geometries, lighting conditions, and lane marking qualities confirm that the proposed look-down system achieves superior real-time performance, comparable lane detection accuracy, and reduced computational overhead relative to both traditional bird's-eye and advanced neural approaches. These findings underscore the practical benefits of a straightforward, explainable, and resource-efficient solution for robust autonomous vehicle lane-keeping.

Keywords: lane detection; autonomous vehicles; PID control; downward-facing cameras; deep learning; bird's-eye view; CARLA simulator; lateral control

check for updates

Academic Editor: Raffaele Di Gregorio

Received: 5 February 2025
Revised: 26 February 2025
Accepted: 3 March 2025
Published: 6 March 2025

Citation: Nadiri, F.; Rad, A.B. A Novel Lateral Control System for Autonomous Vehicles: A Look-Down Strategy. *Machines* **2025**, *13*, 211. https://doi.org/10.3390/machines13030211

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Lane detection and lateral control are among the core functionalities required in modern autonomous vehicles, ensuring roadway safety and reliable navigation [1]. Inspired by human drivers, most lateral control algorithms employ the so-called "look-ahead" strategy [2] whereby a single front-facing camera provides a preview of upcoming lane geometry. Controllers such as Pure Pursuit [3], the Stanley method [4], and model predictive controllers [5] compute steering commands based on the projected vehicle path ahead. These methods generally require a model of the vehicle. By the same token, the current advanced driver-assistance systems (ADAS) and self-driving solutions generally rely on camera sensors to perceive road boundaries, with a front-facing single-camera setup being standard practice. A common approach is to transform the forward-facing view into a bird's-eye view (BEV) [6,7] in order to simplify the lane geometry and facilitate line-fitting algorithms. However, BEV-based methods often struggle in closely spaced lanes or during lane-changing maneuvers, where lateral proximity cues are vital. In this study, by contrast, we explore a 'look-down' strategy that captures lane boundaries directly

beneath the vehicle from two downward-facing cameras, offering distinct advantages for close-proximity lane detection.

In the last few years, deep learning-based methods have surged in popularity. End-to-end networks can learn to map raw images directly to steering commands, while semantic segmentation models specialize in extracting lane lines [8]. Despite substantial performance gains, these methods typically require large, carefully annotated datasets and extensive computational resources. Furthermore, these systems are essentially "black-box" solutions, hindering interpretability—a significant concern for safety-critical systems.

In this paper, we propose a novel two-camera look-down lane detection architecture integrated with a classical PID controller for lateral control. The two side-mounted (driver and the passenger side), downward-facing cameras capture images close to the road surface. This architecture of the proposed system in shown in (Figure 1). The setup enhances the lateral boundary perception and is particularly beneficial for maneuvers in constrained settings such as merging, overtaking, and narrow-lane driving. To elucidate the method's uniqueness in relation to existing lane detection and control approaches, the following key novelties and advantages are highlighted:

1. Infrastructure-free close-proximity sensing: The proposed look-down camera configuration detects lateral boundaries without relying on magnetic markers or embedded lane aids, thereby eliminating the need for infrastructure modifications.
2. Robust lane detection under constrained scenarios: By positioning cameras beneath the side mirrors, the system maintains reliable performance in narrow lanes, situations in which look-ahead single-camera approaches often struggle.
3. Explainable classical control: A classical PID controller demonstrates that, when combined with effective sensor placement, a simpler and more transparent control method can achieve competitive performance compared to complex or black-box deep learning solutions.
4. Comprehensive comparison with existing methods: A thorough experimental evaluation in the CARLA simulator [9] benchmarks the proposed look-down strategy against a traditional bird's-eye view approach and an end-to-end deep learning framework (PilotNet) [8], revealing improved accuracy, lower computational demand, and enhanced real-time capability.

The results validate that the proposed look-down method, despite being simpler than many advanced deep learning techniques, excels in terms of real-time performance, and precise lane-keeping. This study also demonstrates that the conventional control principles remain highly relevant, particularly for resource-constrained or safety-critical applications demanding transparent system logic.

Related Work

Lane detection and lateral control are crucial research areas for autonomous driving. Existing methodologies can be categorized into model-based geometric methods [10], traditional computer vision approaches [11], deep learning frameworks [12], and preview control strategies [13]. A comparative study by Artuñedo et al. (2024) systematically evaluated multiple lateral control strategies, including LQR, MFC, SAMFC, PID, and NLMPC, highlighting the trade-offs in stability and performance under different driving conditions [14]. Each of these categories contributes uniquely to the understanding and advancement of lane detection and lateral control techniques.

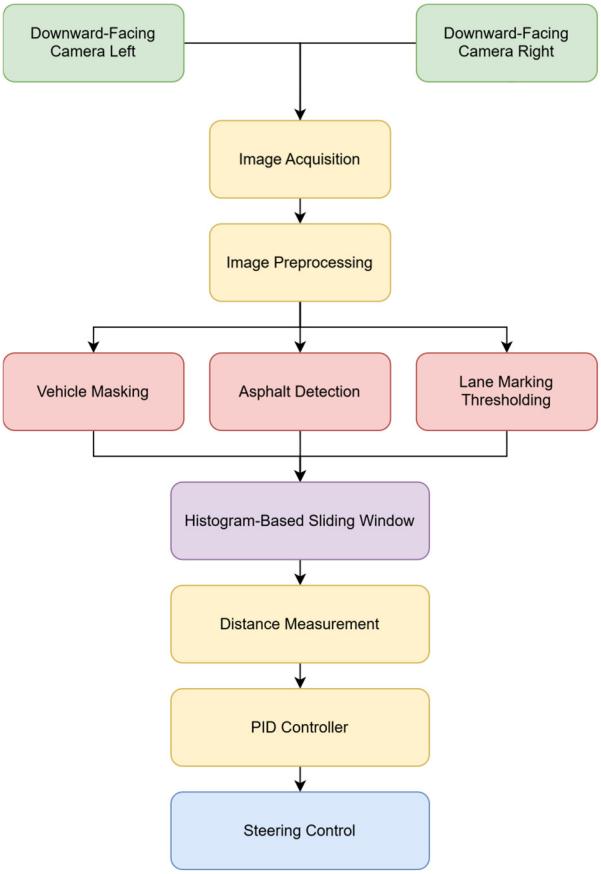

Figure 1. System architecture: Overview of the dual-camera look-down lane detection system integrated with a PID-based lateral controller.

Bird's-eye view (BEV) [15] transformations are widely used in lane detection as they simplify the lane geometry and enable easier extraction of the lane boundaries. Taylor et al. explored vision-based lateral control strategies for highway driving and demonstrated the effectiveness of BEV in structured environments [16]. However, BEV-based methods often struggle with closely spaced lanes, degraded lane markings, and dynamic maneuvers such as lane changes. Similarly, Choi designed a look-down feedback adaptive controller for lateral control in highway vehicles, showing the potential of alternative perspectives like look-down sensing for improved performance in constrained driving conditions [17].

Whereas early studies implemented look-down methods, the literature is rich in look-ahead strategies. Choi pioneered the look-down sensing system [17]. His work focused on discrete magnetic markers rather than visual sensing. Choi demonstrated the utility of adaptive feedback control in structured highway scenarios, where magnetic markers provided reliable lateral position data under various conditions. Although magnetic marker-based systems are robust and cost-effective, they require significant infrastructure modifications, such as embedding markers in the road surface. Li et al. (2021) introduced a preview control approach that adapts steering based on predicted constraint violations, improving vehicle stability under dynamic conditions [18]. This approach aligns with our objective of enhancing lane detection robustness, particularly in constrained environments.

In contrast, this study explores a camera-based look-down approach, which eliminates the need for physical road modifications while retaining the advantages of close-range sensing. The proposed system employs a novel two-camera setup for precise lane boundary detection, particularly in scenarios such as narrow-lane navigation and lane-changing maneuvers. By building on the foundational concept of look-down sensing, this work introduces a modernized, infrastructure-independent alternative using visual data.

Recent advancements in deep learning have significantly enhanced lane detection capabilities. End-to-end approaches and semantic segmentation models have been widely adopted to improve lane detection performance. Chen et al. provided a comprehensive analysis, detailing the advantages of deep learning techniques, including joint feature optimization for perception and planning [19]. Perozzi et al. (2023) proposed a sliding mode control-based lane-keeping system for steer-by-wire vehicles, focusing on smooth transitions between human and autonomous vehicle control [20]. While our method does not incorporate human intervention, it shares the goal of ensuring precise lateral stability in autonomous scenarios. Despite their impressive performance, these methods require substantial computational resources and large annotated datasets. Additionally, their black-box nature raises concerns about interpretability and reliability for safety-critical applications [19,21]. The reliance on data-heavy architectures often makes these methods less viable for real-time deployment in resource-constrained environments. Besides camera-based approaches, sensors like Lidar and radar can provide depth, ranging, and obstacle detection, compensating for camera limitations in poor lighting or inclement weather [22]. Future extensions of the proposed method could integrate these modalities in a sensor-fusion framework, further improving lane detection accuracy and reliability in challenging environments.

Lateral stability remains a critical aspect of autonomous vehicle control. Classical control methods like PID controllers are widely adopted for their simplicity and robust performance in lane-keeping tasks. Advanced control strategies, such as L1 adaptive control, have been developed to handle dynamic uncertainties. Shirazi and Rad applied L1 adaptive control to vehicle lateral dynamics [23], achieving superior robustness under varying conditions. Similarly, Zhu et al. reviewed lateral stability criteria and control applications, emphasizing the importance of robust control frameworks for safety-critical systems [21]. Model Predictive Control (MPC) has also been extensively studied for integrating safety and performance constraints in highway scenarios [24]. However, its computational demands often limit its application in real-time environments.

Comparative studies are essential for understanding the trade-offs between different approaches. Ni et al. conducted an evaluation of MPC-based controllers, demonstrating their ability to address safety constraints in highway driving scenarios [24]. Reinforcement learning-based lateral control methods have also shown promise, as demonstrated by Wasala et al. (2020), where deep reinforcement learning was used to optimize vehicle steering in both simulated and real-world environments [25]. While our approach favors classical PID-based control for its explainability and efficiency, RL-based strategies could offer an adaptive framework for real-time tuning of control parameters. Simulation platforms like CARLA have become standard for validating lane detection and control strategies under diverse conditions [9]. While L1 adaptive control, robust linear parameter-varying controllers, and MPC strategies have proven effective for path tracking, we employ a PID controller due to its simplicity, robustness, low computational overhead, and interpretability. As validated in our experiments (Section 3), properly tuned PID gains can yield state-of-the-art lane-keeping performance under varied conditions without incurring the higher computational costs typical of advanced model-based methods. Under extreme driving conditions or high-speed maneuvers, robust model-based controllers become critical to account for non-linearities and vehicle dynamics [26]. While our study focuses on a resource-friendly PID approach, we acknowledge that advanced methods—such as robust MPC or L1 adaptive control—could be integrated to handle abrupt changes in tire friction or highly non-linear vehicle behavior. In future work, however, we plan to extend our approach by investigating whether techniques such as robust MPC could further improve tracking accuracy and performance in highly dynamic driving scenarios. Geometric path

tracking algorithms (e.g., Pure Pursuit, Stanley) also are favored for their conceptual simplicity and straightforward geometric interpretation. However, they often assume idealized kinematic constraints and may underperform when tire forces, slip angles, or other vehicle dynamics complexities become significant, particularly at high speeds or on slippery road surfaces [27,28]. Consequently, while geometric methods excel in moderate-speed and well-marked scenarios, they typically require careful calibration or more sophisticated dynamic modeling to ensure robust performance under extreme conditions.

This study employs an evaluation framework, comparing the proposed look-down system with BEV and deep learning-based approaches to validate its performance comprehensively.

2. Materials and Methods

This section outlines the proposed dual-camera lane detection and the PID-based lateral control system. Additionally, a comparative study against a single-camera bird's-eye view approach is presented.

2.1. Camera Setup and Data Acquisition

The proposed system leverages three cameras within the CARLA simulation environment to test and validate the proposed method under a variety of conditions (Figure 2). These include:

1. Two downward-facing side cameras: Mounted beneath the driver and the passenger side mirrors, respectively. These cameras capture the road surface directly, enabling precise lane boundary detection.
2. A forward-facing camera: Positioned at the vehicle's front center. This camera provides a full frame of the road.



(**a**) (**b**) (**c**)

Figure 2. Sample frames were captured from the CARLA simulation environment. The frames include: (**a**) a downward-facing view from the left-side camera; (**b**) a forward-facing camera; and (**c**) a downward-facing view from the right-side camera.

The cameras operate at high resolution and real-time frame rates within the simulation environment. In our final 'look-down + PID' approach, only the two downward-facing side cameras supply the lateral boundary information to compute the offset error for the PID controller. The forward-facing camera is included in the simulation scene to (1) allow direct comparisons with conventional front-facing Bird's-Eye View (BEV) methods, and (2) provide input frames for the end-to-end PilotNet baseline. We do not directly fuse the forward-facing data into our primary controller, in order to demonstrate that side-mounted, look-down imagery alone can achieve robust lane-keeping. To evaluate system performance comprehensively, various scenarios were designed in CARLA, including straight and curved roads, lane merges, uphill/downhill sections, and driving under diverse lighting conditions (daylight, dusk, and night) and weather scenarios (clear skies,

rain, and fog). Lane markings with varying qualities—clear, faded, or intermittent—were also simulated to assess robustness.

The simulation provides a controlled environment for recording images and testing lateral control, with each camera capturing synthetic data that is inherently aligned with the ground-truth information, such as lane boundaries and vehicle orientation. This allows for precise analysis and evaluation of the proposed system without reliance on real-world data acquisition.

2.2. Image Preprocessing

A two-stage thresholding process is implemented for each frame to ensure reliable extraction of lane markings:

1. Vehicle Masking: Any visible portion of the vehicle (e.g., wheels, bumpers) is removed via a region-of-interest (ROI) mask. This mitigates the risk of false positives from the vehicle's own edges or reflections.
2. Asphalt Detection with Convex Hull: Color-based thresholding isolates the asphalt region, employing adaptive thresholds to account for shading and texture variations in CARLA. A convex hull refines the asphalt area by enclosing contiguous pixels, eliminating small noisy patches such as debris or shadows.
3. Lane Marking Thresholding: A second color-based threshold highlights potential lane markings (white or yellow) in the asphalt mask. The resulting binary mask contains candidate lane lines, which feed into downstream detection algorithms.

2.3. Lane Detection Methodology

Histogram-Based Sliding Window

A histogram-based approach was adopted for robust lane detection, especially under curved or partially faded markings (Figure 3). The steps are:

1. Histogram Computation: Summation of pixel intensities along the horizontal dimension of the binary mask highlights prominent lane-line locations.
2. Sliding Window Search: Starting from the bottom of the frame, small vertical windows move toward identified histogram peaks, clustering pixels belonging to lane markings.
3. Polynomial Fitting: Aggregated lane points in each window are fitted to a second-order polynomial, capturing potentially curved lane geometry.

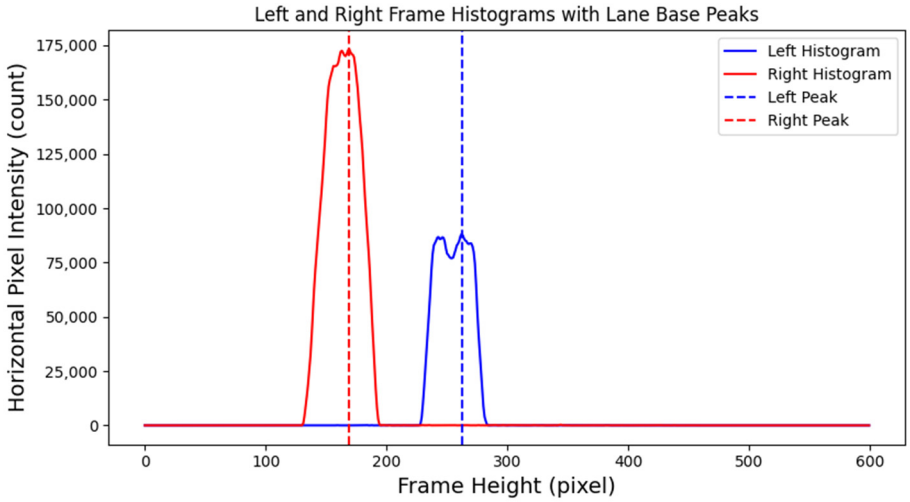

Figure 3. Right and left lane line histogram peaks.

This method provides accurate tracking even on winding roads where simpler line-fitting (e.g., Hough Transform) struggles.

Figure 4 illustrates the sliding window approach applied for lane line detection on a sharp curved road, demonstrating the method's ability to accurately track lane boundaries in challenging road geometries.

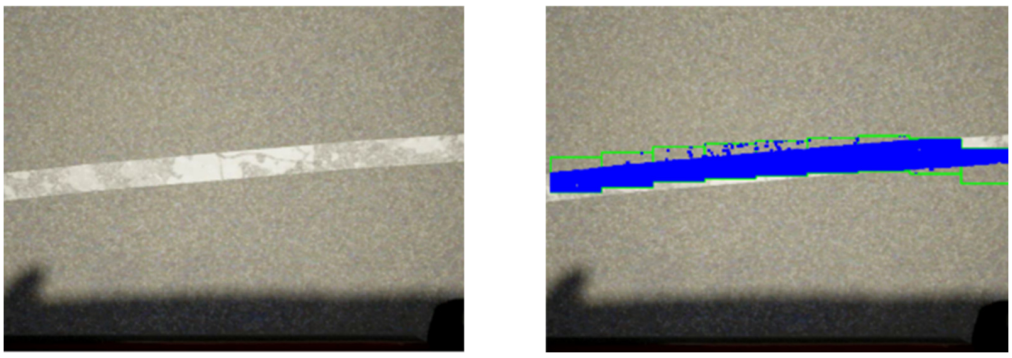

(**a**) (**b**)

Figure 4. Sliding window detection: Illustration of the sliding window approach for lane line detection in a curve road: (**a**) original frame; and (**b**) processed frame with histogram line detection.

2.4. Distance Measurement and Error Calculation

After lane line detection, we compute:

D_R = distance from the vehicle center to the right lane using the right down-facing camera
D_L = distance from the vehicle center to the left lane using the left down-facing camera

An error term:

$$e = D_L - D_R \tag{1}$$

indicates how off-center the vehicle is within the lane. Ideally, $e \approx 0$ implies perfect centering.

2.4.1. PID-Based Lateral Control

A PID controller generates steering commands δ to drive $e \to 0$:

$$\delta(t) = K_P e(t) + K_I \int e(t)dt + K_D \frac{de(t)}{dt} \tag{2}$$

where K_P, K_I, and K_D are the proportional, integral, and derivative gains, respectively. We tune in these gains via simulation-based sweeps, ensuring rapid error correction without excessive overshoot.

A PID controller is used in the method for the following reasons:

1. Computational Simplicity: Minimal overhead, suitable for real-time.
2. Explainability: Transparent control structure, crucial for safety validation.
3. Proven Reliability: Well-established in automotive and industrial control applications.

2.4.2. Controller Parameter Selection

Once the lane boundaries are identified from the thresholded images, we compute the lateral offset $e(t)$ by measuring the pixel distance from the image's centerline. This offset is scaled by the known camera geometry to represent the physical lateral distance in meters. The PID controller then takes $e(t)$ as input to compute the steering angle δ. We performed a parametric sweep over candidate K_p, K_i, and K_d values in the CARLA simulator, aiming to minimize both steady-state offset and overshoot. The final gains were chosen to ensure a rapid response (under 1 s to settle) without oscillatory steering behavior.

Figure 5 illustrates the projected car position after implementing a PID controller for the steering wheel. It demonstrates that the car consistently remained close to position

zero, indicating the car was positioned near the center of the lane. This highlights the effectiveness of the PID controller in maintaining the vehicle's lateral position and keeping it aligned with the lane center, which is crucial for safe and stable lateral control during driving (Figure 6).

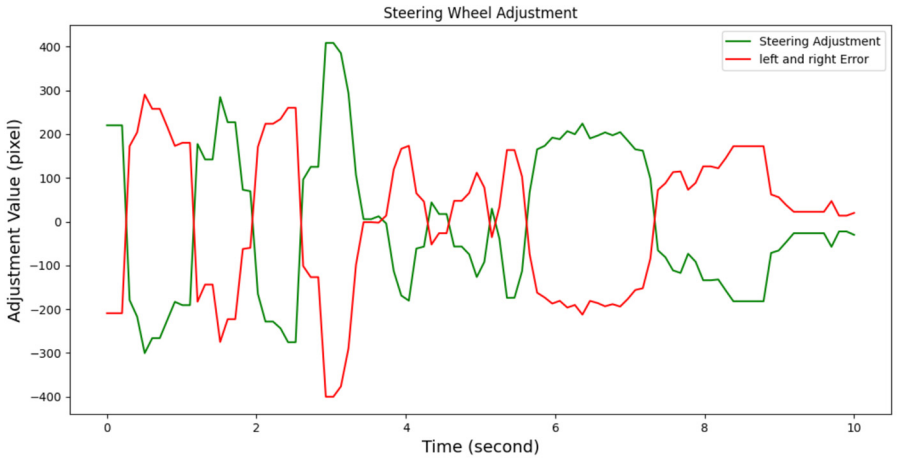

Figure 5. Graph showing the vehicle's lateral error and steering adjustments over time.

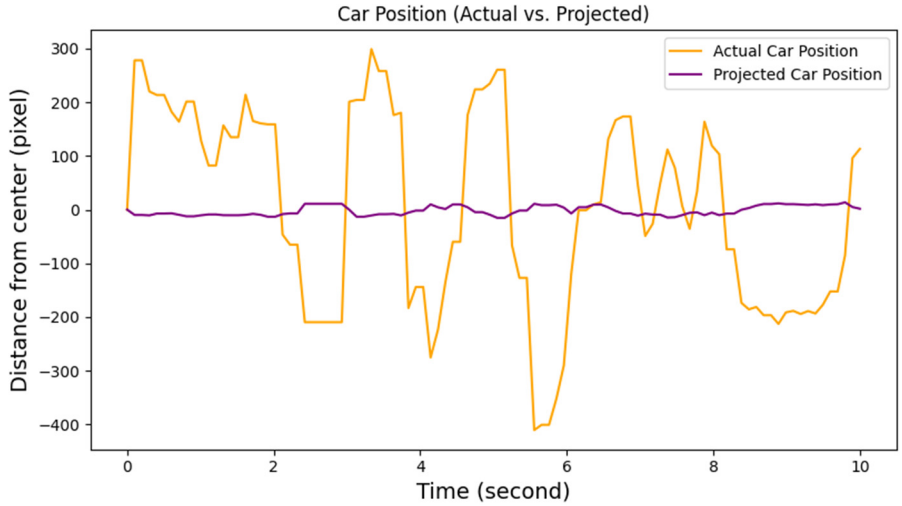

Figure 6. Actual car position compared to projected car position after applying PID controller.

2.5. Experimental Setup

All experiments were conducted on a desktop workstation equipped with an Intel Core i7-12700K CPU (8 cores at 3.6 GHz) manufactured by Intel Corporation, Santa Clara, CA, USA, 32 GB of RAM, and an NVIDIA RTX 3080 GPU (10 GB VRAM) manufactured by NVIDIA Corporation, Santa Clara, CA, USA. CARLA Simulator version 0.9.xx was used on Ubuntu 20.04 with Python 3.8.

We generated data in CARLA by driving the vehicle under various lighting (day/dusk/night) and weather (clear/rain/fog) conditions. We masked out the vehicle hood and wheels by applying a fixed region of interest and performed color thresholding in HSV space to separate lane markings from the asphalt. We then applied morphological operations (closing with a 3×3 kernel) to remove noise.

3. Results

Experiments are conducted in the CARLA simulator, which provides a high-fidelity urban driving environment with configurable roads, traffic, weather, and lighting. We focus on the following scenarios:

1. Straight Roads: Evaluate baseline lane centering accuracy.
2. Curved Roads: Assess performance on windy segments requiring more frequent steering adjustments.
3. Uphill/Downhill: Test robustness to perspective changes.
4. Varied Lighting: Daylight, dusk, and night, capturing changes in contrast and potential glare.
5. Weather: Light rain, heavy rain, and fog.
6. Lane Marking Quality: Sharp, faded, or missing lines.

3.1. Metrics

Measured items:

1. Lane Detection Accuracy: Proportion of frames in which lane boundaries are correctly identified.
2. Lateral Error (RMS): Root Mean Square of $(D_L - D_R)/2$ with respect to the true lane center.
3. Steering Smoothness: Standard deviation of steering commands.
4. Computational Load: Processing time per frame (FPS) and CPU/GPU usage.

3.2. Comparative Methods

3.2.1. Bird's-Eye View and PID Controller

Using the forward-facing camera, we performed a perspective transformation to create a BEV image of the road (see Figure 7). The histogram-based sliding window algorithm is then applied in the transformed space and a PID controller is applied to keep the car in the center of the lane. This forms the traditional baseline for comparison.

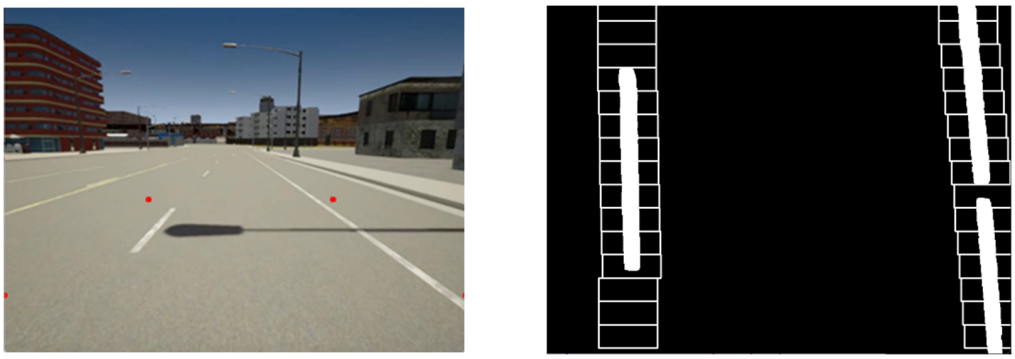

(**a**) (**b**)

Figure 7. Bird's-eye view detection: Transformation of the front-facing camera image to BEV and subsequent lane detection: (**a**) original look-ahead frame with red dots indicating the BEV region of interest (ROI); and (**b**) processed BEV frame with histogram line detection.

3.2.2. Deep Learning Framework

As part of the benchmark for the proposed dual-camera, look-down lane detection, and lateral control framework, the well-known end-to-end driving network PilotNet was selected. Originally developed by Bojarski et al. [8,29], PilotNet replaces traditional feature extraction and rule-based control pipelines with direct mapping from raw image inputs to

steering commands. This design has demonstrated effectiveness when lane markings are well-defined and environmental conditions are relatively stable.

The PilotNet implementation in this study follows the design specifications provided by Bojarski et al. [8]. However, for use in the CARLA simulator, the architecture has been adapted to operate on 160 × 120 input images and produce three control outputs: steering, throttle, and brake [30]. After incoming camera frames are resized to 160 × 120 pixels, they are processed by five convolutional layers. The first three layers use 5 × 5 filters with strides of 2 (24, 36, and 48 output channels), and the remaining two layers use 3 × 3 filters (stride of 1) with 64 output channels. The feature maps are subsequently flattened and passed through four fully connected layers of sizes 1152, 100, 50, and 10 neurons. Each fully connected layer employs ReLU activation, with a dropout rate of 0.1 to mitigate overfitting. Finally, three parallel output nodes (steering, throttle, brake) each compute a mean-squared error (MSE) loss during training. Figure 8 illustrates a high-level schematic of the PilotNet architecture used in these experiments.

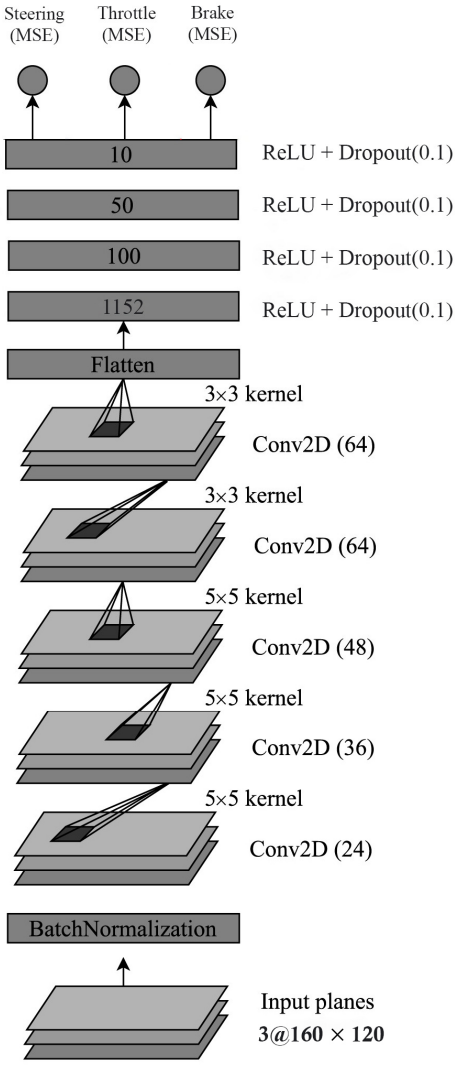

Figure 8. High-level schematic of the PilotNet architecture used in the experiments.

PilotNet was not pretrained in this study. Instead, a supervised end-to-end training strategy was adopted using data collected within the CARLA simulator. A single forward-facing camera was attached to a vehicle under autopilot, and each image frame was labeled with the ground-truth steering, throttle, and brake signals from the simulator. Several thousand images were recorded across a variety of map layouts and lighting/weather condi-

tions, then partitioned into approximately 75% for training and 25% for validation/testing. Shuffling was performed to reduce any sequential correlation in the data. All images were normalized to a $[0, 1]$ range prior to model input.

An Adam optimizer with a learning rate of 1×10^{-4} and a batch size of 64 was used. Training typically lasted for 30 epochs, with early stopping based on the validation loss. The overall loss is defined as:

$$Loss_{total} = MSE\left(\theta_{pred}, \theta_{gt}\right) + MSE\left(t_{pred}, t_{gt}\right) + MSE\left(b_{pred}, b_{gt}\right) \quad (3)$$

where θ, t, and b denote steering angle, throttle, and brake, respectively.

PilotNet's end-to-end structure generally provided robust lane-following performance under consistent lighting and clearly visible lane markings. The method benefited from direct image-to-steering mapping, removing the need for explicit hand-engineered features. However, it also exhibited limitations in scenarios with degraded or missing lane markers, highlighting its dependence on discernible visual cues. In addition, PilotNet's black-box nature reduced interpretability, making it challenging to trace specific decision factors within the learned representations. Nevertheless, these attributes make PilotNet an informative baseline to compare against more specialized or sensor-fusion-based approaches.

In the subsequent sections, PilotNet's performance metrics and behavior under varied conditions are contrasted with those of the proposed dual-camera system. This comparison examines trade-offs related to robustness, accuracy, and interpretability, providing insights into the respective merits of end-to-end deep learning frameworks versus more structured lane detection and control pipelines.

3.3. Comparative Analysis of Lane Detection and Lateral Control Methods

To comprehensively evaluate the performance of the proposed dual-camera look-down lane detection and PID-based lateral control system, we conducted a series of experiments in the CARLA simulator. The system was benchmarked against the conventional bird's-eye view (BEV) + PID controller [7] and PilotNet. The evaluation focused on detection accuracy, lateral error, computational efficiency, and steering smoothness. The following sections present a consolidated comparison supported by a single table for clarity.

Table 1 highlights that the look-down + PID system offers a balanced performance with high detection accuracy, low lateral error, excellent steering smoothness, and superior computational efficiency compared to both traditional BEV methods and advanced deep learning frameworks.

Table 1. Consolidates the key quantitative metrics for an overall comparison of the different lane detection and lateral control methods.

Metric	Look-Down + PID	BEV + PID	PilotNet
Lane Detection Accuracy (%)	≈95	≈91	≈90
Average Lateral Error (Straight Roads, m)	<0.07	≈0.12	≈0.20
Average Lateral Error (Curves, m)	<0.12	≈0.18	≈0.25
Steering Smoothness (Std. Dev., rad)	≈0.05	≈0.08	≈0.12
Frame Rate (FPS)	≈40	≈30	≈22
CPU Utilization (%)	≈35	≈40	≈50
GPU Utilization (%)	<10	≈40	≈70

Table 1 clearly shows that while deep learning models like PilotNet achieve good lane detection accuracy under different conditions, they do so with significantly higher

computational demands. The look-down + PID method, on the other hand, offers a robust alternative with excellent real-time performance and lower resource requirements.

3.3.1. Detection Accuracy and Stability

The look-down method consistently achieves high accuracy (≈95%) in detecting lane lines, particularly excelling in curved or congested scenarios. This additional lateral perspective proves beneficial during lane changes, outperforming the BEV approach and handling partial markings more gracefully. In contrast, the BEV + PID method attains good performance on straightforward roads but degrades during lane changes and fails more frequently when lane markings are significantly faded or occluded. Deep learning framework PilotNet maintains moderate accuracy (≈90%) and struggles with lane ambiguity and degraded markings, especially in scenarios not represented in the training data.

3.3.2. Lateral Error and Control

The lateral error analysis in Figure 9 demonstrates the PID controller integrated with the look-down setup maintains an average lateral error of <0.07 m on straight roads and <0.12 m on sharp curves, indicating precise lateral positioning and consistent lane centering. The BEV + PID approach shows larger error variances, particularly during quick lane-change maneuvers. The deep learning-based PilotNet framework has 0.20 m on straight roads and 0.25 m on curves.

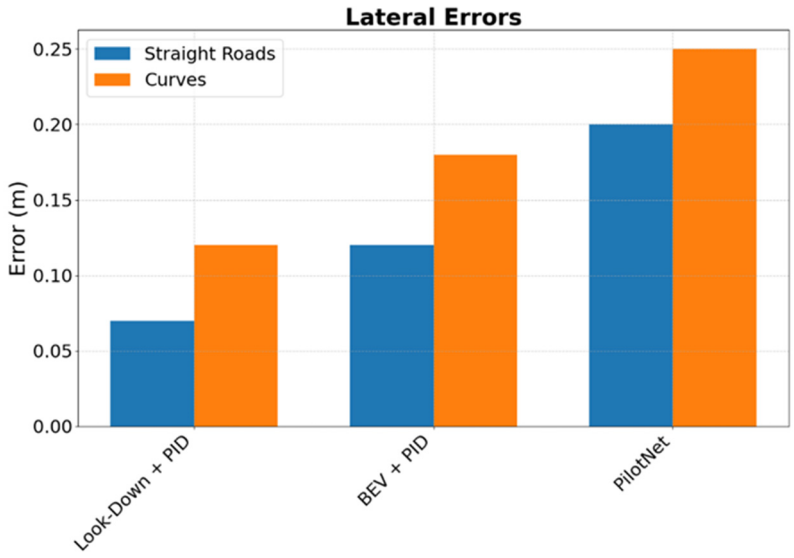

Figure 9. Lateral error and control comparison.

While deep learning models can achieve near-comparable RMS errors to the look-down system, they do so at a significantly higher computational cost.

3.3.3. Computational Efficiency

The system performance metrics in Figure 10 highlight significant advantages in computational efficiency:

1. Processing Speed: The look-down + PID system maintains ≈40 FPS, substantially outperforming both BEV + PID (≈30 FPS) and PilotNet (≈22 FPS). This higher frame rate enables more responsive control and better adaptation to dynamic scenarios.
2. Resource Utilization:
    - CPU usage remains modest at ≈35% for look-down + PID, compared to ≈40% for BEV + PID and ≈50% for PilotNet.

- GPU utilization shows even more dramatic differences: <10% for look-down + PID versus ≈40% for BEV + PID and ≈70% for PilotNet.

3. Scalability Implications: The lower resource requirements of the system make it more suitable for deployment on cost-effective hardware platforms while maintaining superior performance metrics.

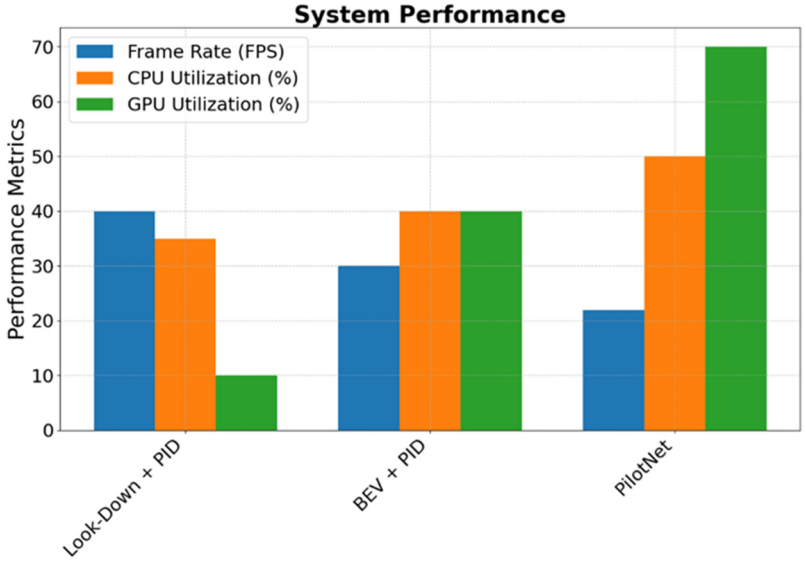

Figure 10. System performance metrics comparison.

3.4. Qualitative Observations

Beyond quantitative metrics, qualitative assessments provide deeper insights into the system's performance under various real-world conditions. The look-down method consistently maintains clear lane boundary detection even in challenging scenarios such as heavy rain and night-time glare, whereas the BEV method shows reduced clarity under the same conditions. Additionally, steering angle trajectories during dynamic maneuvers reveal that the look-down + PID system achieves smoother and more precise steering adjustments, ensuring stable lane-keeping.

3.5. Longitudinal Speed Comparison

Although our main focus is on lateral control, we also recorded the longitudinal speed profiles for each method—(i) look-down + PID, (ii) look-ahead + PID, and (iii) PilotNet—to investigate how speed varies during lane-keeping tasks. Unlike the first two methods, which apply a fixed (or minimally adaptive) throttle, PilotNet inherently learns throttle and brake from training data, potentially allowing it to slow down on curves.

3.5.1. Experimental Setup for Speed

1. Look-down + PID/Look-ahead + PID:
   - Steering is controlled by a PID loop.
   - Throttle is set to a constant 10 m/s (36 km/h) on straight roads. On curved roads, speed is reduced to 8 m/s (28.8 km/h) to improve lane stability and minimize lateral deviation, unless otherwise specified.
   - No active braking is employed.
2. PilotNet (End-to-End):
   - Learns both steering and throttle/brake from the same forward-facing camera feed.
   - Tends to decelerate on sharper curves and accelerate on straights.

In our implementation of the look-down + PID approach, we employ a simple rule-based (open-loop) throttle schedule rather than an adaptive longitudinal controller. Specifically, we fix the throttle to maintain a speed of approximately 10 m/s on straight segments of the road and manually reduce it to around 8 m/s whenever the vehicle enters a curved section. This reduction is triggered by a basic road curvature threshold—no active feedback or real-time speed adaptation is involved. By contrast, PilotNet modulates speed dynamically, decelerating more aggressively on sharper curves and accelerating on straights based on its learned end-to-end policy.

3.5.2. Speed Profiles in Curved and Straight Segments

We evaluated the methods on two representative routes in CARLA:

- Moderately curved roads
- S-Curve with a short straight section

Table 2 presents the average speed, standard deviation, and min/max speeds observed for each method. Figure 11 shows the longitudinal speed vs. time on curved roads.

Table 2. Speed statistics for moderately curved roads.

Method	Longitudinal Control	Avg. Speed (m/s)	Std. Dev. (m/s)	Min/Max (m/s)
Look-Down + PID	~Fixed at 10 m/s	≈9.8	±0.8	7.8/10.2
Look-Ahead + PID	~Fixed at 10 m/s	≈10.1	±0.7	8.1/10.5
PilotNet	Learned Throttle and Brake	≈8.5	±1.0	7.0/10.2

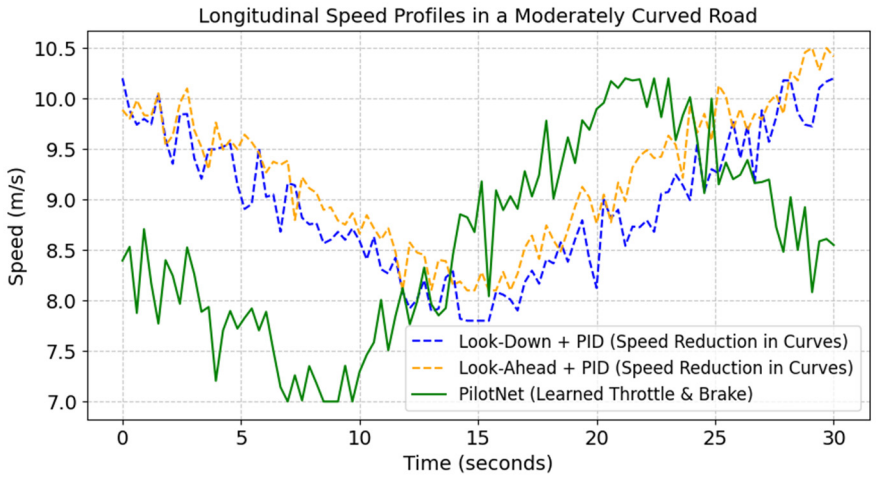

Figure 11. Speed vs. time on a curved road. Look-down and look-ahead PID controllers have constant speed, while PilotNet dynamically adjusts throttle and braking.

Key Observations:

- Look-down + PID and look-ahead + PID adjust their speed, reducing to ≈8 m/s in curved sections to improve lane stability and minimize lateral deviation. This leads to a wider speed range (min: 7.8 m/s for look-down + PID and 8.1 m/s for look-ahead + PID), showing controlled deceleration in curves.
- PilotNet inherently learns throttle and braking, showing more aggressive deceleration in sharper curves (min: 7.0 m/s) but with a higher speed variance (±1.0 m/s) due to its adaptive control approach.

4. Discussion

The findings in this study underscore that classical control strategies can remain competitive alongside modern deep learning methods, especially when combined with an optimal camera configuration. While deep learning excels in perfect conditions and can generalize diverse scenarios after adequate training, it risks failure if domain mismatches occur or if it encounters environmental conditions that are not present in the training set. Moreover, its high computational and data labeling demands can be prohibitive in real-time, mass-deployment scenarios.

In contrast, the look-down + PID approach:

- Requires minimal tuning and computational resources.
- Maintains interpretability and predictable failure modes.
- Excels in close-proximity detections, a scenario sometimes overlooked by front-facing or top-down cameras.

Although our proposed system focuses on lateral control with a near-constant speed, the results in Section 3.5 illustrate how coupling speed adaptation can improve cornering stability. The look-ahead and look-down PID methods do not inherently slow on curves unless a predefined speed schedule is introduced. By contrast, PilotNet automatically modulates the throttle, showing lower average speeds on sharper curves. However, this approach demands extensive training and is less transparent. Future work will integrate a more advanced longitudinal controller (e.g., MPC or an adaptive rule-based approach) with the look-down method to optimize both speed and steering. To further clarify the comparative advantages of different lateral control methods, Table 3 summarizes key performance aspects [14,31–34].

Table 3. Comparative advantages of different lateral control methods.

Method	Complexity	Vehicle and Road Model	Robustness	Performance in Curved or Uphill/Downhill Roads
Pure pursuit	Moderate	Required	Low	Poor
Stanley	Moderate	Required	Moderate	Moderate
Model Predictive	High	Required	High	Moderate to High
Deep Learning	High (extensive training is required)	Model free	Moderate	Moderate
Proposed method	Low	Model free	High	High

The comparison highlights that while deep learning-based approaches and model predictive controllers can provide improved performance under specific conditions, the proposed look-down PID method offers a computationally efficient and robust alternative, particularly in challenging environments like sharp curves and uneven terrain.

Future research could explore a hybrid approach that integrates semantic segmentation to enhance lane detection in complex environments, particularly in scenarios with occlusions, varying lighting conditions, or degraded lane markings. The use of vision-based transformers for lane detection could significantly improve feature extraction and lane representation, enabling more robust performance in challenging conditions. These models could be trained to capture long-range dependencies in road scenes, leading to better generalization across diverse environments.

For lateral control, a simpler Proportional-Integral-Derivative (PID) controller may suffice for standard driving scenarios, while a more advanced Model Predictive Control (MPC) approach could be employed for improved trajectory optimization and smoother steering adjustments, especially in dynamic road conditions. The combination of deep

learning-based lane detection with classical control strategies could strike a balance between accuracy and computational efficiency.

Additionally, sensor fusion techniques incorporating LiDAR or radar could further enhance robustness under adverse weather conditions such as heavy rain, fog, or snow. These sensors can provide complementary information to mitigate the limitations of camera-based lane detection, particularly in cases where lane markings are poorly visible.

Future implementations could also investigate reinforcement learning-based controllers that dynamically adjust driving policies based on real-time environmental feedback, leading to adaptive and optimal lane-keeping performance. Moreover, integrating multi-modal sensor inputs with an attention-based fusion mechanism could refine localization and control accuracy, pushing autonomous lane-keeping closer to human-like adaptability.

5. Conclusions

This paper presents a two-camera look-down system for lane detection and PID-based lateral control in autonomous vehicles. By mounting cameras beneath the side mirrors and capturing the road surface, this method significantly improves detection accuracy and steering stability during lane changes and tight maneuvers. Comprehensive comparisons against both a bird's-eye view baseline and advanced deep learning methods validate the system's real-time efficiency, robust detection performance, and strong lateral control while imposing modest hardware requirements.

Key takeaways include:

1. Enhanced Proximity Awareness: Downward-facing cameras excel at capturing lateral boundaries even in constrained scenarios.
2. Computational Simplicity: The PID controller and localized image regions enable high processing speeds on conventional hardware.
3. Robustness and Interpretability: Clear, predictable behavior compared to computationally intensive or black-box deep learning solutions.
4. Viable Real-World Application: Suitable for scenarios where reliability, low cost, and transparency are paramount.

In summary, the dual-camera, look-down approach achieves a 95% lane detection accuracy, maintains average lateral errors below 0.07 m on straight roads and 0.12 m in tight curves, and processes camera inputs at 40 FPS with less than 10% GPU usage. These results underscore the real-time feasibility and robustness of the proposed system for practical autonomous driving applications. Future investigations may focus on extending the system to multi-lane highways with heavy traffic, lane changing, integrating obstacle detection, and leveraging adaptive or learning-based controllers in a hybrid manner. Overall, the new look-down lane detection and lateral control strategy provides a practical, efficient, and highly interpretable pathway toward safer and more dependable autonomous vehicles.

Author Contributions: Conceptualization, A.B.R.; methodology, A.B.R. and F.N.; software, F.N.; validation, A.B.R. and F.N.; formal analysis, A.B.R.; investigation, A.B.R.; resources, A.B.R.; data curation, F.N.; writing—original draft preparation, F.N.; writing—review and editing, A.B.R.; visualization, F.N.; supervision, A.B.R.; project administration, A.B.R. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: No new external datasets were created or analyzed in this study. The data used in this research were generated within the CARLA simulator for experimental purposes and are available upon request.

Conflicts of Interest: Author Farzad Nadiri was employed by the company Quartech. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

BEV	Bird's-eye view
PID	Proportional-Integral-Derivative
MPC	Model Predictive Control
ROI	Region of Interest
MSE	Mean-squared error
ADAS	Advanced driver-assistance systems

References

1. Ho, M.L.; Chan, P.T.; Rad, A.B.; Shirazi, M.; Cina, M. A novel fused neural network controller for lateral control of autonomous vehicles. *Appl. Soft Comput.* **2012**, *12*, 3514–3525. [CrossRef]
2. Cina, M.; Rad, A.B. We steer where we stare. *Transp. Res. Interdiscip. Perspect.* **2024**, *25*, 101092. [CrossRef]
3. Park, M.-W.; Lee, S.-W.; Han, W.-Y. Development of lateral control system for autonomous vehicle based on adaptive pure pursuit algorithm. In Proceedings of the 2014 14th International Conference on Control, Automation and Systems (ICCAS 2014), Goyang-si, Republic of Korea, 22–25 October 2014; pp. 1443–1447. [CrossRef]
4. Hoffmann, G.M.; Tomlin, C.J.; Montemerlo, M.; Thrun, S. Autonomous Automobile Trajectory Tracking for Off-Road Driving: Controller Design, Experimental Validation and Racing. In Proceedings of the 2007 American Control Conference, New York, NY, USA, 11–13 July 2007; pp. 2296–2301. [CrossRef]
5. Wu, X.; Lin, L.; Chen, J.; Du, S. A Fast Predictive Control Method for Vehicle Path Tracking Based on a Recurrent Neural Network. *IEEE Access* **2024**, *12*, 141104–141115. [CrossRef]
6. Ren, H.; Wang, M.; Lei, X.; Zhang, M.; Li, W.; Liu, C. Refinement Bird's Eye View Feature for 3D Lane Detection with Dual-Branch View Transformation Module. In Proceedings of the ICASSP 2024—2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Seoul, Republic of Korea, 14–19 April 2024; pp. 5930–5934. [CrossRef]
7. Batista, M.P.; Shinzato, P.Y.; Wolf, D.F.; Gomes, D. Lane Detection and Estimation using Perspective Image. In Proceedings of the 2014 Joint Conference on Robotics: SBR-LARS Robotics Symposium and Robocontrol, Sao Carlos, Brazil, 18–23 October 2014; pp. 25–30. [CrossRef]
8. Bojarski, M.; Del Testa, D.; Dworakowski, D.; Firner, B.; Flepp, B.; Goyal, P.; Jackel, L.D.; Monfort, M.; Muller, U.; Zhang, J.; et al. End to End Learning for Self-Driving Cars. *arXiv* **2016**, arXiv:1604.07316.
9. Dosovitskiy, A.; Ros, G.; Codevilla, F.; Lopez, A.; Koltun, V. CARLA: An Open Urban Driving Simulator. *arXiv* **2017**, arXiv:1711.03938.
10. Lombard, A.; Buisson, J.; Abbas-Turki, A.; Galland, S.; Koukam, A. Curvature-Based Geometric Approach for the Lateral Control of Autonomous Cars. *J. Franklin Inst.* **2020**, *357*, 9378–9398. [CrossRef]
11. Cao, J.; Song, C.; Song, S.; Xiao, F.; Peng, S. Lane Detection Algorithm for Intelligent Vehicles in Complex Road Conditions and Dynamic Environments. *Sensors* **2019**, *19*, 3166. [CrossRef]
12. Kocić, J.; Jovičić, N.; Drndarević, V. An End-to-End Deep Neural Network for Autonomous Driving Designed for Embedded Automotive Platforms. *Sensors* **2019**, *19*, 2064. [CrossRef]
13. Zhu, G.; Hong, W. Autonomous Vehicle Motion Control Considering Path Preview with Adaptive Tire Cornering Stiffness Under High-Speed Conditions. *World Electr. Veh. J.* **2024**, *15*, 580. [CrossRef]
14. Artuñedo, A.; Moreno-Gonzalez, M.; Villagra, J. Lateral control for autonomous vehicles: A comparative evaluation. *Annu. Rev. Control* **2024**, *57*, 100910. [CrossRef]
15. Can, Y.B.; Liniger, A.; Unal, O.; Paudel, D.; Van Gool, L. Understanding Bird's-Eye View of Road Semantics using an Onboard Camera. *IEEE Robot. Autom. Lett.* **2022**, *7*, 3302–3309. [CrossRef]
16. Taylor, C.J.; Košecká, J.; Blasi, R.; Malik, J. A Comparative Study of Vision-Based Lateral Control Strategies for Autonomous Highway Driving. *Int. J. Rob. Res.* **1999**, *18*, 442–453. [CrossRef]
17. Choi, S.B. The design of a look-down feedback adaptive controller for the lateral control of front-wheel-steering autonomous highway vehicles. *IEEE Trans. Veh. Technol.* **2000**, *49*, 2257–2269. [CrossRef]
18. Li, R.; Ouyang, Q.; Cui, Y.; Jin, Y. Preview control with dynamic constraints for autonomous vehicles. *Sensors* **2021**, *21*, 5155. [CrossRef]

19. Chen, L.; Wu, P.; Chitta, K.; Jaeger, B.; Geiger, A.; Li, H. End-to-end Autonomous Driving: Challenges and Frontiers. *IEEE Trans. Pattern Anal. Mach. Intell.* **2024**, *46*, 10164–10183. [CrossRef]
20. Perozzi, G.; Rath, J.J.; Sentouh, C.; Floris, J.; Popieul, J.C. Lateral Shared Sliding Mode Control for Lane Keeping Assist System in Steer-by-Wire Vehicles: Theory and Experiments. *IEEE Trans. Intell. Veh.* **2023**, *8*, 3073–3082. [CrossRef]
21. Zhu, Z.; Tang, X.; Qin, Y.; Huang, Y.; Hashemi, E. A Survey of Lateral Stability Criterion and Control Application for Autonomous Vehicles. *IEEE Trans. Intell. Transp. Syst.* **2023**, *24*, 10382–10399. [CrossRef]
22. Viadero-Monasterio, F.; Alonso-Rentería, L.; Pérez-Oria, J.; Viadero-Rueda, F. Radar-Based Pedestrian and Vehicle Detection and Identification for Driving Assistance. *Vehicles* **2024**, *6*, 1185–1199. [CrossRef]
23. Shirazi, M.M.; Rad, A.B. L1 Adaptive Control of Vehicle Lateral Dynamics. *IEEE Trans. Intell. Veh.* **2018**, *3*, 92–101. [CrossRef]
24. Ni, L.; Gupta, A.; Falcone, P.; Johannesson, L. Vehicle Lateral Motion Control with Performance and Safety Guarantees. *IFAC-PapersOnLine* **2016**, *49*, 285–290. [CrossRef]
25. Wasala, A.; Byrne, D.; Miesbauer, P.; O'Hanlon, J.; Heraty, P.; Barry, P. Trajectory based lateral control: A Reinforcement Learning case study. *Eng. Appl. Artif. Intell.* **2020**, *94*, 103799. [CrossRef]
26. Viadero-Monasterio, F.; Nguyen, A.-T.; Lauber, J.; Boada, M.J.L.; Boada, B.L. Event-Triggered Robust Path Tracking Control Considering Roll Stability Under Network-Induced Delays for Autonomous Vehicles. *IEEE Trans. Intell. Transp. Syst.* **2023**, *24*, 14743–14756. [CrossRef]
27. Qiang, W.; Yu, W.; Xu, Q.; Xie, H. Study on Robust Path-Tracking Control for an Unmanned Articulated Road Roller Under Low-Adhesion Conditions. *Electronics* **2025**, *14*, 383. [CrossRef]
28. Zhang, J.; Yang, M.; Chen, Y.; Li, H.; Chen, N.; Zhang, Y. Research on sliding mode control for path following in variable speed of unmanned vehicle. *Optim. Eng.* **2025**. [CrossRef]
29. Bojarski, M.; Yeres, P.; Choromanska, A.; Choromanski, K.; Firner, B.; Jackel, L.; Muller, U. Explaining How a Deep Neural Network Trained with End-to-End Learning Steers a Car. *arXiv* **2017**, arXiv:1704.07911.
30. Paniego, S.; Shinohara, E.; Cañas, J. Autonomous driving in traffic with end-to-end vision-based deep learning. *Neurocomputing* **2024**, *594*, 127874. [CrossRef]
31. Lee, Y.; You, B. Comparison and Evaluation of Various Lateral Controller for Autonomous Vehicle. In Proceedings of the 2020 International Conference on Electronics, Information, and Communication (ICEIC), Barcelona, Spain, 19–22 January 2020; pp. 1–4. [CrossRef]
32. Ayyagari, A.; Ahmad, R. Comparative Study of Model-Based Lateral Controllers with Selected Deep Learning Methods for Autonomous Driving. Master's Thesis, Simon Fraser University, Vancouver, BC, Canada, 2023.
33. Tiwari, T.; Agarwal, S.; Etar, A. Controller Design for Autonomous Vehicle. In Proceedings of the 2021 International Conference on Advances in Electrical, Computing, Communication and Sustainable Technologies (ICAECT), Bhilai, India, 19–20 February 2021; pp. 1–5. [CrossRef]
34. Vivek, K.; Sheta, M.A.; Gumtapure, V. A Comparative Study of Stanley, LQR and MPC Controllers for Path Tracking Application (ADAS/AD). In Proceedings of the 2019 IEEE International Conference on Intelligent Systems and Green Technology (ICISGT), Visakhapatnam, India, 29–30 June 2019; pp. 67–674. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.