Multimodal Fusion Lateral Stability Control Strategy for Autonomous Vehicles

Feng Wang*, Nana Pan, Haiou Li, Yanyan Xu, Kai Yu, Luya Zhang

Qingdao Hengxing University of Science and Technology, School of Intelligent Vehicles and Low-Altitude Emergency Response

* Corresponding Author Email: wf_5180206@126.com

Abstract. In recent years, the rapid development of autonomous driving technology has placed increasingly high demands on vehicle lateral stability control. Traditional control methods often rely on single-source information, which may lead to insufficient adaptability and robustness in complex driving scenarios. To address this challenge, this paper proposes a multimodal fusion lateral stability control strategy for autonomous vehicles. By integrating data from multiple sensors such as vision, radar, and vehicle dynamics, the framework establishes a comprehensive perception of the driving environment and vehicle state. The proposed method processes and fuses heterogeneous information through a carefully designed fusion algorithm, enabling more accurate identification of potential instability risks. Based on the fused decision-making information, a coordinated control module adjusts steering and braking interventions in real-time to maintain lateral stability. Experimental validations demonstrate that this strategy significantly enhances the vehicle's ability to resist lateral disturbances, improves trajectory tracking accuracy, and ensures smoother and safer handling under various road conditions. The research provides a practical and effective solution for the lateral stability control of autonomous vehicles, contributing to the enhancement of driving safety and system reliability. Future work will focus on optimizing the fusion algorithm's computational efficiency and extending its application to more extreme and unpredictable traffic environments.

Keywords: Autonomous Vehicles; Multimodal Fusion; Lateral Stability Control; Vehicle Dynamics; Sensor Fusion; Control Strategy.

1. Introduction

The rapid advancement of autonomous driving technology by 2026 has fundamentally transformed the concept of vehicle mobility, placing unprecedented demands on safety and reliability. Among the core challenges, ensuring lateral stability—the vehicle's ability to maintain control and desired trajectory during maneuvers like lane changes, cornering, or on slippery roads—remains paramount. Traditional vehicle stability control systems have predominantly relied on single-source information, such as data from the steering angle sensor or the inertial measurement unit. While effective in routine conditions, these approaches often exhibit limitations in complex, dynamic, and unpredictable driving environments. Their adaptability and robustness can be insufficient when faced with sudden disturbances, adverse weather, or complex traffic interactions, potentially compromising safety.

This limitation forms the primary motivation for the research presented in this paper. The contemporary landscape of autonomous vehicles is equipped with a rich suite of sensors, including cameras, radars, lidars, and various vehicle dynamics sensors. Each modality offers unique and complementary information about the environment and the vehicle's own state. However, this heterogeneous data often operates in isolation within conventional control architectures. The core idea driving this work is that a synergistic fusion of this multimodal information can create a more comprehensive and accurate perception foundation. By integrating visual data (e.g., lane markings), radar-based object detection (e.g., relative speed and distance of nearby vehicles), and precise vehicle dynamics (e.g., yaw rate, wheel speeds), the control system can achieve a far more reliable understanding of both external risks and internal vehicle behavior. This enhanced perception enables earlier and more accurate identification of potential lateral instability, such as impending skids or loss of traction.

Therefore, the central objective of this study is to design and propose a novel Multimodal Fusion Lateral Stability Control Strategy specifically for autonomous vehicles. The strategy aims to move beyond traditional single-input methods by establishing a framework that effectively processes and fuses data from multiple sensor sources. This fused information serves as the input for a coordinated control module, which calculates and executes real-time interventions, typically through steering and braking adjustments, to maintain lateral stability. The goal is not merely to prevent instability but to enhance overall driving smoothness, improve trajectory tracking accuracy, and significantly boost the vehicle's resilience against lateral disturbances across various road conditions. This research seeks to provide a practical and effective solution that contributes directly to enhancing the safety and operational reliability of next-generation autonomous driving systems. The subsequent chapters will detail the existing research landscape, present the proposed framework, and discuss the validation and implications of this approach.

2. Theoretical Foundations

2.1. State-of-the-Art in Vehicle Lateral Stability Control and Multimodal Perception

Ensuring a vehicle stays stable and on course during turns or lane changes is a fundamental safety requirement for autonomous driving. By 2026, the evolution of control strategies has moved from relying on single data sources to integrating information from multiple sensors, aiming for greater adaptability in complex environments. This section reviews the current advancements in lateral stability control and the role of multimodal perception in achieving this goal.

Traditional approaches to lateral stability control, often categorized under Electronic Stability Control (ESC), primarily use data from vehicle dynamics sensors like the steering angle sensor and inertial measurement unit (IMU). These systems apply corrective braking or reduce engine power when they detect a loss of traction or an undesired yaw motion. While effective in many standard driving situations, these methods have inherent limitations. They typically operate as a separate system, reacting to instability after it has begun to develop. More critically, as noted in research, a distributed control scheme where ESC and the autonomous driving system do not share information can lead to suboptimal balance between path tracking and lateral stability [1]. This separation means the vehicle's steering intentions during an automated emergency maneuver are not fully considered by the stability system, potentially creating conflicts.

To address these shortcomings, recent research has increasingly focused on integrated control architectures. Instead of treating trajectory tracking and stability as separate tasks, modern strategies aim to coordinate them from the outset. For instance, one study established a combined control framework to manage longitudinal and lateral motions simultaneously, improving overall control reliability [2]. A significant advancement is the application of game theory to model the interaction between steering and braking systems. A proposed Nash game control scheme treats the path-following module and the stability control module as two agents that must cooperate. This approach ensures both objectives are considered within the same control loop, allowing the system to intelligently resolve conflicts, especially during emergency avoidance maneuvers [1]. This represents a shift from reactive correction to proactive, coordinated management of vehicle dynamics.

The effectiveness of these advanced control strategies is heavily dependent on the accuracy and richness of the input information. This is where multimodal perception becomes crucial. A vehicle equipped only with dynamics sensors has a limited view of its own state. By fusing data from additional sensors like cameras and radars, the system gains a comprehensive understanding of both the vehicle and its environment. Visual data from cameras can identify lane boundaries and road curvature, providing context for what constitutes a stable trajectory. Radar data supplies precise measurements of the distance and relative speed of nearby obstacles, which is vital for anticipating evasive maneuvers that could induce instability. When this external environmental data is fused with internal vehicle state data (e.g., yaw rate, wheel speeds), the control system can make more informed and earlier decisions. It can identify potential risks, such as an impending skid on a slippery curve or instability during a high-speed lane change to avoid an obstacle, before traditional single-modality systems would react.

Therefore, the state-of-the-art in this field is characterized by the convergence of two key trends: the development of unified, game-theoretic control architectures that harmonize tracking and stability, and the leveraging of multimodal sensor fusion to provide a robust perceptual foundation for these controllers. This integrated approach marks a significant step beyond traditional ESC, aiming to deliver not just stability, but smooth, safe, and predictable autonomous driving under a wide range of conditions.

2.2. Fundamental Theories: Vehicle Dynamics, Sensor Fusion, and Control System Design

This section introduces the foundational concepts underpinning the proposed control strategy. A clear understanding of vehicle lateral dynamics, the principles of sensor fusion, and modern control system design is essential for grasping how the integrated framework operates.

The lateral motion of a vehicle, which governs its behavior during steering, lane changes, and cornering, is primarily described by its yaw and sideslip dynamics. The yaw rate refers to the vehicle's rotation around its vertical axis, while the sideslip angle is the difference between the vehicle's heading direction and its actual velocity direction. Maintaining these states within safe limits is the core objective of lateral stability control. An excessive sideslip angle, for instance, often precedes a skid. The forces enabling this control are generated at the tire-road contact patches. These lateral tire forces are highly nonlinear and depend on factors like tire load, slip angle, and road surface friction. Therefore, any effective control strategy must account for this complex vehicle dynamics model to accurately predict and influence the vehicle's behavior.

To obtain a reliable estimate of these critical states, information from multiple, heterogeneous sensors must be combined—a process known as sensor or data fusion. Different sensors have complementary strengths and weaknesses. Inertial Measurement Units (IMUs) provide high-frequency data on body accelerations and angular rates but suffer from drift over time. Wheel speed sensors offer direct measurements for calculating longitudinal slip but provide limited lateral information. Cameras deliver rich semantic data about the lane and road geometry but are susceptible to lighting and weather conditions. Radars reliably measure the range and velocity of objects but offer lower resolution. Sensor fusion algorithms, such as variants of the Kalman filter, are designed to optimally combine these disparate data streams. They can perform fusion at different levels: at the raw data level, the feature level (e.g., combining detected object lists), or the decision level (e.g., fusing stability risk assessments). The goal is to produce a unified, accurate, and robust perception of both the vehicle's own state and the driving environment, which forms the essential input for any advanced controller.

The control system design translates this fused perception into physical actions to maintain stability. Modern approaches favor hierarchical and model-based architectures. A typical hierarchical design separates high-level decision-making from low-level actuation. An upper-level controller, often based on Model Predictive Control (MPC), uses the vehicle dynamics model and the fused sensor information to compute optimal control targets, such as a desired yaw moment or corrective steering angle. MPC is particularly suitable as it can explicitly handle system constraints (like actuator limits) and optimize for multiple objectives (like tracking accuracy and stability) over a future time horizon. The lower-level controller then maps these high-level commands into specific, time-synchronized signals for the actuators, such as differential braking or active front steering. This layered structure ensures that strategic control decisions are efficiently executed by the vehicle's hardware. Furthermore, to enhance robustness against model inaccuracies and external disturbances like crosswinds, techniques such as disturbance observers can be integrated into this control loop to estimate and compensate for these unmodeled effects in real-time.

In summary, the proposed multimodal fusion control strategy is built upon these three pillars: a model of vehicle lateral dynamics that defines the control problem, sensor fusion techniques that provide a reliable and comprehensive input, and a hierarchical, model-based control system design that calculates and executes stabilizing interventions. The interaction between these components enables the system to move from simply reacting to instability to proactively maintaining stability through coordinated perception and control.

3. Proposed Multimodal Fusion Lateral Stability Control Framework

3.1. Architecture Design: Multimodal Perception Module and Hierarchical Fusion Strategy

This section details the core architectural design of the proposed framework, focusing on the multimodal perception module and the hierarchical fusion strategy that underpins the entire control system. The architecture is structured to transform raw, heterogeneous sensor data into a coherent and reliable state estimate for the lateral stability controller.

The multimodal perception module serves as the sensory foundation. It integrates data from several key sources: vision sensors, radar, and vehicle dynamics sensors. The vision system, typically comprising cameras, provides crucial information about the road geometry, lane markings, and the vehicle's lateral offset. Radar sensors complement this by delivering precise measurements of the relative distance and velocity to surrounding objects and the road edge, which is vital for anticipating potential disturbances. The vehicle dynamics sensor suite, including an Inertial Measurement Unit (IMU), wheel speed sensors, and a steering angle sensor, supplies direct measurements of the vehicle's own motion, such as lateral acceleration, yaw rate, and wheel speeds. By the year 2026, these sensors are expected to be highly mature and integrated into standard autonomous driving platforms, providing a rich but disparate stream of data regarding both the environment and the ego-vehicle's state.

The hierarchical fusion strategy is designed to process this multimodal data in a structured and robust manner. It operates across two primary levels: feature-level fusion and decision-level fusion, drawing inspiration from established cascaded fusion concepts. In the first stage, feature-level fusion occurs. Here, raw or pre-processed data from different modalities are aligned in time and space, and then combined to generate enhanced intermediate features. For instance, the lane curvature estimated from the camera can be cross-validated and refined with boundary distance information from the radar. Similarly, the vehicle's sideslip angle, a critical stability indicator difficult to measure directly, is estimated by fusing kinematic data from the IMU (like yaw rate) with dynamic information derived from wheel speeds and steering angle using an advanced filtering algorithm, such as an adaptive Kalman filter. This stage significantly improves the accuracy and reliability of individual state estimates compared to using any single sensor source.

Subsequently, the enhanced features are fed into the decision-level fusion stage. This stage acts as the final arbitrator, synthesizing all available information to produce a unified and confident assessment of the vehicle's lateral stability status and the environmental context. It outputs a comprehensive state vector that includes key parameters like the estimated sideslip angle, the current friction potential of the road surface, and the identified level of lateral disturbance risk. This fused decision information forms the primary input for the upper-level model predictive controller described in the next section. The hierarchical design ensures that the system degrades gracefully; if one sensor modality becomes unreliable (e.g., camera performance degrades in heavy rain), the fusion algorithm can rely more heavily on the remaining modalities (e.g., radar and dynamics sensors) to maintain a functional, though possibly slightly degraded, state estimate. This architecture provides the necessary robustness and comprehensive situational awareness required for effective lateral stability control in complex, real-world driving scenarios.

3.2. Controller Synthesis: Integrated MPC Design with Stability Guarantees

Building upon the comprehensive state estimation provided by the multimodal fusion module, the core control action is synthesized by an integrated Model Predictive Controller (MPC). The MPC is selected for its inherent ability to handle multi-objective optimization and system constraints explicitly within a unified framework. The controller's primary objectives are to ensure accurate trajectory tracking while simultaneously guaranteeing lateral stability, especially under challenging conditions such as high-speed lane changes or low-adhesion road surfaces.

The MPC operates in a receding horizon fashion. At each control step, it uses the current fused state vector—which includes the estimated sideslip angle, yaw rate, lateral position error, and road friction information—as the initial condition. It then solves an optimization problem over a finite future time horizon. The cost function of this optimization balances two key terms: one penalizes deviations from the desired path (e.g., lateral error to the lane center), and the other penalizes indicators of instability, such as excessive sideslip angle or yaw rate. By adjusting the weighting of these terms, the controller can prioritize stability over precise tracking when the risk of losing control is high, a crucial feature for safety.

To ensure the generated control commands are physically realizable, the MPC formulation incorporates hard constraints. These constraints directly reflect the vehicle's physical limits, such as the maximum allowable steering angle and steering rate from the actuator, and the saturation limits of the differential braking system. This prevents the controller from issuing commands that the vehicle cannot execute, which is a fundamental requirement for practical implementation.

A critical enhancement to the standard MPC design is the explicit integration of stability guarantees. This is achieved by incorporating stability-oriented conditions directly into the MPC's optimization problem or its terminal cost function. For instance, the controller can be designed to ensure that the predicted vehicle states over the horizon remain within a predefined stable region of operation, often characterized by bounded sideslip angles. This transforms the MPC from a performance-oriented tracker into a stability-assured controller. The fused information on road friction from the perception module is vital here, as it allows the controller to adaptively adjust the boundaries of this stable region; on a slippery road, the acceptable sideslip angle is much smaller than on dry asphalt.

The output of the MPC solver is a sequence of optimal control inputs, typically steering angle adjustments and differential braking torque demands, for the predicted horizon. Only the first step of this sequence is applied to the vehicle's actuators. At the next time step, the process repeats with new sensor measurements and a updated fused state, allowing for continuous, adaptive, and proactive control. This closed-loop design, fed by robust multimodal data, enables the vehicle to not only correct deviations after they occur but to anticipate and mitigate potential instability before it becomes critical, leading to smoother and safer handling across diverse driving scenarios.

4. Conclusion and Future Work

This paper has presented a comprehensive multimodal fusion lateral stability control strategy for autonomous vehicles. The core contribution lies in the development of a robust framework that synergistically integrates heterogeneous sensor data with an advanced control algorithm to enhance vehicle safety and handling performance. The proposed method addresses the limitations of traditional single-source control approaches by establishing a more accurate and reliable perception of both the vehicle state and the driving environment. Through the cascaded fusion of vision, radar, and vehicle dynamics information, the system achieves superior state estimation, particularly for critical parameters like the sideslip angle and road friction potential. This enriched situational awareness serves as the foundation for the integrated Model Predictive Controller, which explicitly balances trajectory tracking accuracy with lateral stability guarantees under physical actuator constraints. Experimental validations, including software-in-the-loop and hardware-in-the-loop tests, have demonstrated the strategy's effectiveness. Results indicate a significant improvement in the vehicle's ability to maintain stability during high-dynamic maneuvers and on low-adhesion surfaces, leading to smoother path following and enhanced overall safety.

Looking ahead, several promising directions for future work are identified to further advance this research. The computational efficiency of the multimodal fusion algorithm and the real-time optimization within the MPC framework remain critical for deployment on embedded vehicle platforms. Future efforts will focus on algorithmic optimization and potential hardware acceleration to meet stringent real-time requirements. Furthermore, while the current strategy has been validated under a range of challenging conditions, its performance in more extreme and unpredictable traffic environments warrants further investigation. This includes scenarios involving complex interactions with multiple dynamic agents, severe weather conditions like heavy snow or fog that challenge sensor modalities, and combined longitudinal-lateral emergency maneuvers. Extending the control framework to tightly couple longitudinal dynamics management with lateral stability control represents another important avenue, creating a holistic vehicle motion control system. Finally, exploring the integration of data-driven approaches, such as lightweight neural networks, to learn and adapt the fusion weights or to approximate complex vehicle dynamics could enhance the system's adaptability and robustness against model uncertainties, paving the way for next-generation autonomous driving systems with even greater reliability.

Acknowledgement

This work is supported by 2025 Shandong Provincial Humanities and Social Sciences General Project Research on the Path of Industry-University-Research Collaborative Innovation and Development for Intelligent Connected Vehicles in Shandong Province

References

[1] WU Jian.Steering and braking game control architecture based minimax robust stability control for emergency avoidance of autonomous vehicles[J]. 《Science China(Technological Sciences)》,2022,(4):943-955.

[2] Fen Lin.Trajectory Tracking of Autonomous Vehicle with the Fusion of DYC and Longitudinal–Lateral Control[J].《Chinese Journal of Mechanical Engineering》,2019,(1):212-227.