J Intel Connect Veh

Short Communication

Physics-informed machine learning for autonomous driving: From perception to control

Yubing Yang[1], Meng Li[2], Yuhao Wang[3,✉]

Cite this article: Yang Y B, Li M, Wang Y H. *J Intel Connect Veh* 2026, **9**(1): 9210077. https://doi.org/10.26599/JICV.2026.9210077

1 Introduction

With the rapid development of autonomous driving (AD), achieving safe, efficient, and interpretable decision-making in complex traffic environments remains a major challenge for AD. Currently, AD faces several real-world challenges, including handling observation noise, occlusions, and sensor blind spots. Furthermore, data-driven methods are constrained by the scarcity of training data for long-tail or rare scenarios and the high costs of data collection and annotation. These factors diminish the physical consistency and generalization capabilities of data-driven methods. Additionally, planning and control in multiagent interactions can easily yield infeasible or unsafe decisions, making it challenging to meet the stringent safety and compliance requirements of the real world. Introducing physical knowledge into data-driven methods, utilizing physics as structured prior information, can improve learning efficiency and reduce collection and annotation costs when data are scarce (Banerjee et al., 2025). Crucially, it provides dynamic constraints and safety constraints during the decision-making process, thereby increasing the model's interpretability and verifiability and laying the foundation for building safe, reliable, and scalable AD systems.

In this context, physically informed machine learning (PIML) offers a promising solution. The core idea of PIML is to systematically embed prior knowledge of the physical world (from macroscopic conservation laws and symmetries to microscopic vehicle dynamics equations) into a data-driven learning process. This method not only implements physical equations as regularization terms in the loss function but also encompasses various technical approaches, including designing model structures with physical constraints, constructing hybrid models that fuse physics and data, and performing physics-guided feature learning (Banerjee et al., 2025; Karniadakis et al., 2021). However, existing work focuses more on macroscopic levels. There remains a research gap regarding how PIML can be leveraged within the perception, planning, and control modules to achieve physical consistency, risk awareness, and verifiable safe control.

This study aims to discuss the significance and application potential of PIML within key AD modules. We specifically examine its potential utilization and limitations in three critical areas: physically consistent state estimation, trajectory prediction based on physical priors and safety constraints, and adaptive and verifiable control strategies. We regard PIML as a modular enhancement method. By embedding physical constraints and residual regularization into the perception, planning, and control modules, we can improve reliability and interpretability under conditions of noise, incomplete observations, or sample scarcity (Fig. 1). However, achieving formal safety guarantees and real-time deployment also requires integration with control theoretic verification methods such as invariant sets, control barrier function (CBF), and robust model predictive control, as well as consideration of engineering-level implementation. Based on this, this study is structured around the following three scientific questions:

1) How to reconstruct physically consistent and spatiotemporal continuous traffic scenarios under incomplete observation;

2) How to achieve safe trajectory prediction for multiple traffic participants under physical constraints;

3) How to achieve verifiable, safe, and adaptive control in partially known environments.

2 Perception

2.1 Physically consistent local state estimation

In complex traffic environments, sensor noise, occlusion, and incomplete or abnormal observations can lead to discontinuities or abrupt changes in state estimation. To achieve robust state estimation, kinematic and dynamic constraints can be introduced during the detection and tracking stages to suppress noise and incomplete observation information. The core idea is to leverage the laws of motion as prior constraints for state estimation. In the temporal dimension, the deviation of state variables (such as position, velocity, and acceleration) from fundamental motion equations can be embedded as a regularization term in the loss function using automatic differentiation or constrained optimization, thereby enforcing temporal smoothness and physical feasibility between adjacent frames. In the spatial dimension, dynamic consistency and geometric constraints are utilized to achieve continuous trajectory completion, preventing state discontinuities caused by occlusion or sparse sampling. Introducing physical consistency into state estimation strengthens

[1] School of Information Engineering, Chang'an University, Xi'an 710064, China. [2] School of Mechanical and Aerospace Engineering, Nanyang Technological University, Singapore 639798, Singapore. [3] School of Vehicle and Mobility, Tsinghua University, Beijing 100084, China.

✉ Corresponding author. E-mail: yuhao6@tsinghua.edu.cn

Received: December 15, 2025; Revised: January 4, 2026; Accepted: January 20, 2026

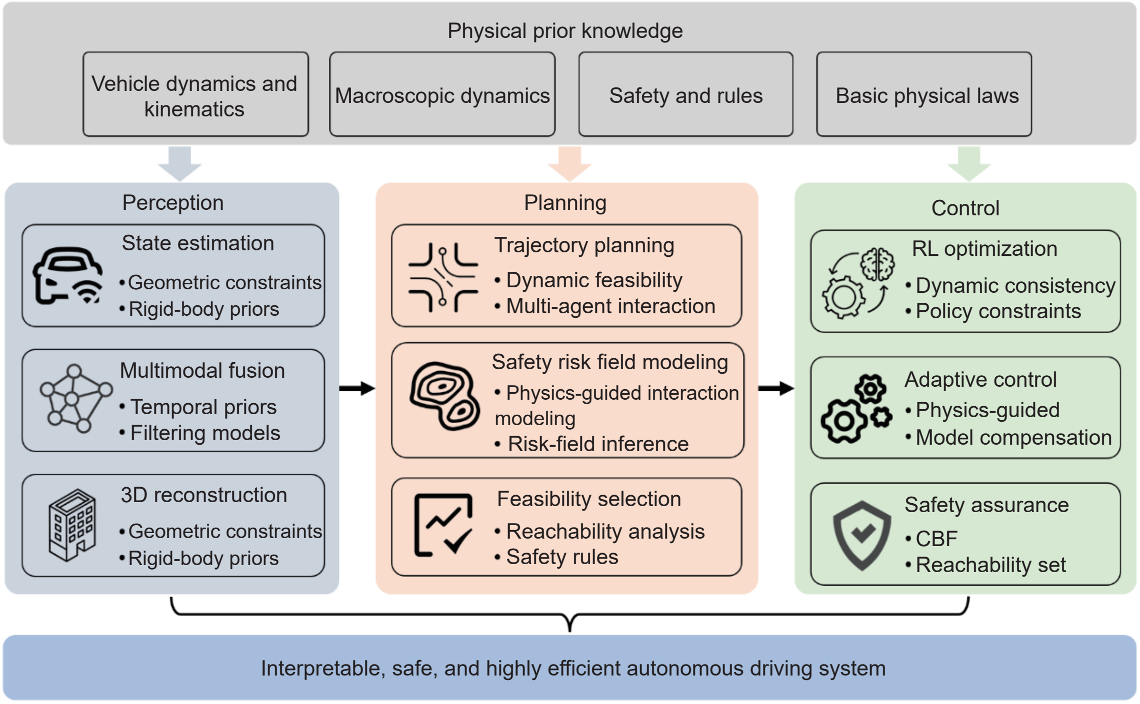

Fig. 1 Framework of physics-informed machine learning for an enhanced autonomous driving system.

the perception stage and enables more robust state estimation. It effectively mitigates the noise effect, compensates for incomplete observations, and preserves spatiotemporal continuity. As a result, the downstream planning and control modules receive more accurate and reliable inputs.

2.2 Spatio-temporally consistent multimodal fusion

Building on local state estimation, further spatiotemporally consistent fusion is crucial for enhancing the robustness and continuity of perception. Sensors (such as cameras, millimeter-wave radar, and LiDAR) differ in sampling frequency, resolution, and latency. Traditional fusion methods based on time alignment or weighted averaging often lead to issues of state desynchronization and semantic inconsistency. To address this, a synchronization mechanism based on kinematic constraints can be introduced in the temporal dimension to guarantee physical feasibility across observation frames. In the spatial dimension, geometric consistency constraints can complete multimodal registration. During fusion, observations that violate kinematic or geometric constraints are not simply rejected; they are downweighted or corrected through physics-constrained filtering, where the motion model provides a prior for smoothing and compensating inconsistent samples. Furthermore, during the fusion stage, physics-constrained filtering is employed, using vehicle or pedestrian motion models as priors for state evolution, to achieve synergistic alignment and trajectory optimization of cross-modal information in the spatiotemporal domain (Long et al., 2024).

2.3 Physically interpretable 3-dimensional (3D) structure reconstruction

In the 3D geometric reconstruction stage, introducing mechanics and rigid body constraints ensures that the perception results are structurally physically reasonable, dynamically stable, and interpretable. Embedding constraints such as contact dynamics, cross-view geometric consistency, or energy conservation conditions into point cloud registration, depth estimation, or scene reconstruction networks can effectively mitigate errors caused by occlusion and sparse observations. This physically interpretable 3D reconstruction process provides high-fidelity environmental input for the planning and control modules, thereby supporting physically consistent planning and control.

3 Planning

3.1 Physical consistency of trajectory prediction

In complex traffic environments, the planning module must generate physically feasible trajectories under multiagent interactions and dynamic constraints. To ensure the safety and feasibility of the planned trajectory, vehicle dynamics constraints can be embedded into the trajectory generation model, such as a nonholonomic dynamics model and higher-order dynamic equations. By incorporating dynamic residual regularization into the prediction loss using PIML, the velocity, acceleration, and lateral offset of the generated trajectory are made temporally continuous and have satisfactory dynamic feasibility (Zhang et al., 2024). This mechanism not only mitigates unreasonable predictions generated by traditional data-driven models in long-tail or sparse sample scenarios but also significantly improves the model's generalization ability and physical interpretability.

3.2 Multiagent interaction and safety risk field modeling

In scenarios with multiple traffic participants, trajectory planning not only needs to satisfy individual vehicle dynamics but also considers the nonlinear coupling among participants. Therefore, interaction-aware mechanisms inspired by potential fields or social force models can be incorporated into the trajectory prediction model to capture the interactions among traffic participants, thereby ensuring collision avoidance, safe distances, and compliance with traffic rules. Furthermore, we introduce the safety risk field to serve as a bridge between perception and planning. The risk field is a spatiotemporal continuous scalar field whose value is mapped by the states of all surrounding traffic participants, the ego-vehicle state, and the road structure. Based on the distribution of the risk field, different candidate trajectories are evaluated and optimized to generate a trajectory with minimum risk. Ultimately, this risk field can be used to guide the trajectory generation process, enabling active collision avoidance, maintaining safe distances, and establishing a physically consistent link between macroscopic situational awareness and microscopic trajectories.

3.3 Rapid feasibility and safety screening

To balance real-time performance and safety, the generated trajectory should undergo rapid physical consistency and safety feasibility verification. Currently, generative methods that directly output trajectories tend to produce “hallucinations” and lack sufficient reliability. Therefore, PIML can be used to integrate physical priors into the trajectory selection process. The planning module generates multiple candidate trajectories from perception data, which are then rapidly evaluated and filtered according to physical constraints such as vehicle dynamics, traffic rules, and safe distances. This mechanism functions as a physical feasibility filter, enhancing the reliability of trajectory planning while maintaining real-time performance. It also reduces the computational burden for subsequent optimization, decision-making, and control, providing highly feasible inputs.

4 Control

4.1 Reinforcement learning (RL) policy optimization via physical residual regularization

In dynamic or unknown environments, RL-based control policies may produce actions that violate physical consistency or dynamic feasibility. To address this, a physical residual regularization term can be incorporated into the RL training process. This integrates vehicle dynamics constraints or energy conservation as soft constraints, forcing the policy update to follow physical priors during action generation. For instance, by minimizing the residual between the rate of change of acceleration, yaw rate, or tire slip and the dynamic model, the generation of unstable or infeasible actions can be reduced. This method achieves physical consistency constraint learning at the policy level, improving the interpretability and reliability of RL policies in real-world traffic environments.

4.2 Real-time parameter estimation and adaptive control based on PIML

Traffic environments are subject to significant uncertainties (such as road friction coefficients, slopes, vehicle loads, and pedestrian reaction characteristics). Using PIML, a neural network can be constructed to simultaneously learn the structural residuals of known models and the unknown environmental parameters. This creates a nonlinear compensator that is more effective than traditional adaptive control or disturbance observers, allowing the controller to be dynamically adjusted. This adaptive mechanism, which combines physical priors with data-driven learning, enables the controller to maintain robustness against sensor noise, model drift, or unknown scenarios. Consequently, it ensures the handling and safety of the vehicle in complex dynamic environments (Li et al., 2025).

4.3 Safety boundary approximation and verifiable constraints based on a physics-informed neural network (PINN)

To achieve safety-verifiable control, PINNs can be used to approximate a safety boundary function, which is then embedded into the RL or optimal control process to constrain the action space directly. In terms of implementation, this can be combined with CBF or reachability set analysis, merging the PINN-approximated safety boundary with formal safety assurance mechanisms. By introducing CBF conditions or reachability constraints into the loss function, the actions generated by the policy are guaranteed to remain within the physically feasible domain (Tayal et al., 2025). This enables verifiable safe control within the PIML framework. By integrating real-time parameter estimation and physical constraints, this method provides formal safety guarantees for the control policy, allowing the AD system to maintain system safety and stability in unknown environments.

5 Conclusions and prospects

This study discusses the potential application of PIML in addressing core challenges of AD systems: physically consistent state estimation, multitraffic participant trajectory planning, and safe and verifiable control in unknown environments. These studies demonstrate that PIML can improve the robustness and physical consistency of systems in data-scarce and long-tail scenarios and enhance the interpretability of each module.

However, future research still faces two main challenges: At the methodological level, it requires achieving coordinated transmission of physical constraints across the entire process, unified modeling of risk fields and decision optimization, and exploring efficient strategies for embedding physical priors in high-dimensional multimodal environments to balance theoretical rigor, system feasibility, and verifiability. At the engineering level, processes such as multimodal constraint solving, physical canonical gradient propagation, and trajectory selection may increase computational complexity and real-time performance pressure. Therefore, methods such as model compression, low-rank approximation, incremental learning, or online optimization are needed to improve real-time performance and scalability while maintaining physical consistency and interpretability, thus promoting practical applications in vehicular environments.

Overall, PIML provides a verifiable physical consistency framework for AD systems and offers new ideas for physical enhancement and interpretability optimization of end-to-end AD, driving the development of safe, efficient, and reliable ADs in complex traffic environments.

Acknowledgements

This work is sponsored by the National Natural Science Foundation of China (No. 52402386).

Declaration of competing interest

The authors have no competing interests to declare that are relevant to the content of this article.

References

Banerjee, C., Nguyen, K., Fookes, C., George, K., 2025. Physics-informed computer vision: A review and perspectives. *ACM Comput Surv*, **57**, 1–38.

Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., Yang, L., 2021. Physics-informed machine learning. *Nat Rev Phys*, **3**, 422–440.

Li, X., Wang, Y., Ozbay, K., Jiang, Z. P., 2025. Physics-informed machine learning with heuristic feedback control layer for autonomous vehicle control. In: 2025 IEEE Intelligent Vehicles Symposium (IV), 2304–2309.

Long, K., Shi, X., Li, X., 2024. Physics-informed neural network for cross-dynamics vehicle trajectory stitching. *Transp Res Part E Logist Transp Rev*, **192**, 103799.

Tayal, M., Singh, A., Kolathaya, S., Bansal, S., 2025. A Physics-informed Machine Learning Framework for Safe and Optimal Control of Autonomous Systems. https://arxiv.org/abs/2502.11057

Zhang, S., Zhang, J., Yang, L., Chen, F., Li, S., Gao, Z., 2024. Physics guided deep learning-based model for short-term origin–destination demand prediction in urban rail transit systems under pandemic. *Engineering*, **41**, 276–296.

Yubing Yang is currently pursuing the M.S. degree in the School of Information Engineering at Chang'an University, China. His research interests include autonomous driving decision planning, deep reinforcement learning, and intelligent transportation systems.

Meng Li received the Ph.D. degree in transportation engineering from Southeast University, China, in 2024. He is currently a Research Fellow with the School of Mechanical and Aerospace Engineering, Nanyang Technological University, Singapore. His research interests include autonomous driving, traffic flow theory, and intelligent transportation systems.

Yuhao Wang received the M.S. degree in transport from Imperial College London in 2019 and the Ph.D. degree from The Hong Kong Polytechnic University in 2023, focusing on intelligent transport system. He is now a Shuimu Tsinghua Scholar and Postdoctoral Researcher Fellow with the Emerging Transportation Solution Lab, School of Vehicle and Mobility, Tsinghua University, China. His current research interests include artificial intelligence and internet of things, urban aerial mobility, and low-altitude digital twin simulation.