HybridLoss – An Adaptive Planning-Oriented Loss Function for End-to-End Autonomous Vehicle

Donglei Rong, Chengcheng Yang, Congcong Bai, Wentong Guo, Sheng Jin, Min Xu

***Abstract*—Autonomous driving often suffer from a decoupled feedback loop between prediction and planning. While prediction losses focus on the accuracy of surrounding agents, planning losses typically imitate recorded Autonomous Vehicles' (AVs) trajectories that may contain suboptimal or aggressive behaviors, leading to unstable interactions in mixed traffic. This paper presents HybridLoss, an adaptive planning-oriented objective that unifies prediction and planning through planner-in-the-loop supervision and interaction-aware consistency. HybridLoss integrates an adaptive motion-planning module which replaces ground-truth targets with optimized reference trajectories, and a multi-term loss combining prediction, adaptive planning, safety potential, and social force objectives. Evaluations on the INTERACTION dataset indicate that HybridLoss significantly outperforms strong baselines. Beyond standard metric improvements—reducing ADE/FDE from 1.36/1.64 m to 1.11/1.36 m and collision rates from 0.19% to 0.11%—extensive stress-testing reveals superior system maturity. First, HybridLoss exhibits the highest robustness under input perturbations, maintaining the lowest planning deviation and endpoint standard deviation. Second, it demonstrates strong generalization, maintaining stable success rates (87.7%) in unseen scenarios with high computational efficiency (64.3 Hz). Third, multi-objective analysis confirms that HybridLoss achieves the optimal Pareto trade-off between efficiency, safety, and comfort, avoiding the speed-safety collapse seen in baseline methods. Finally, social force evaluations highlight that HybridLoss fosters implicit cooperation, achieving higher yield rates and reduced conflict indices while maintaining safe interaction buffers. These results validate HybridLoss as a robust, socially compliant, and adaptive solution for end-to-end driving.**

***Note to Practitioners*—HybridLoss addresses the challenge of balancing individual safety and global efficiency in AV planning within mixed-traffic environments. HybridLoss dynamically integrates adaptive motion-planning with a multi-component loss function to create closed-loop feedback between prediction accuracy and planning quality. This adaptability allows AVs to rapidly optimize trajectories for diverse scenarios, such as prioritizing collision avoidance for merging vehicles while maintaining platoon efficiency on highways. Practitioners can deploy HybridLoss in urban or highway settings to enhance robustness and E2E performance, though environment-specific tuning of loss weights is recommended. While introducing complexity in unpredictable edge cases, HybridLoss offers a scalable, computation-efficient solution that outperforms conventional approaches by harmonizing individual and collective AV objectives.**

***Index Terms*—HybridLoss, adaptive motion planning, prediction and planning, autonomous vehicle, end-to-end**

I. Introduction

Autonomous Vehicles (AVs) are increasingly deployed under national-level initiatives and industrial roadmaps, improve traffic safety, and accelerate large-scale commercialization. Despite rapid progress, Human-driven Vehicles (HVs) will remain dominant for the foreseeable future. As a result, AVs must operate in mixed-traffic where they frequently interact with HVs in safety-critical maneuvers such as merging, lane changing, unprotected turns, and intersection negotiation [1].

While AVs typically excel at perception, collision avoidance, and rule compliance, interaction with HVs is inherently strategic and uncertain. Drivers may behave aggressively, hesitate, or deviate from "nominal" rules, and their intentions are only partially observable. Learning-based driving stacks often decompose the problem into prediction for surrounding agents and AVs' planning. However, a limitation is that prediction and planning are not tightly coupled during training, and the loss functions used to supervise them do not always translate into stable behaviors in closed-loop. A standard integrated End-to-End (E2E) is typically formulated as the weighted sum of prediction loss and planning loss [2]. While widely used, this objective can inherit undesirable behaviors from demonstrations and does not explicitly ensure that prediction–planning combinations are interaction-consistent, which is critical for mixed-traffic stability.

This gap is illustrated by contrasting two common training paradigms. Prediction-only objectives focus on minimizing trajectory errors of HVs but do not directly constrain the quality of AVs' behavior, leaving the planner susceptible to compounding errors and interaction failures. Integrated prediction–planning frameworks mitigate this by jointly optimizing prediction and planning losses. However, planning

This study was supported by the National Key R&D Program of China (2023YFB4302600), the Funds for International Cooperation and Exchange of the National Natural Science Foundation of China (72361137006), the Research Grants Council of the Hong Kong Special Administrative Region, China (PolyU 15224824) and the Research Committee of The Hong Kong Polytechnic University (4-ZZSF). (*Corresponding author: Sheng Jin*)

Donglei Rong is with the Institute of Intelligent Transportation Systems, College of Civil Engineering and Architecture, Zhejiang University, Hangzhou 310058, China, and also with the Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Hong Kong, China (e-mail: donglei_rong@zju.edu.cn)

Chengcheng Yang is with the School of Vehicle and Mobility, Tsinghua University, Beijing 100084, China (e-mail: yang-cc@tsinghua.edu.cn)

Congcong Bai, and Wentong Guo are with the Institute of Intelligent Transportation Systems, College of Civil Engineering and Architecture, Zhejiang University, Hangzhou 310058, China (e-mail: 12112125@zju.edu.cn, 22260182@zju.edu.cn)

Sheng Jin is with the Institute of Intelligent Transportation Systems, College of Civil Engineering and Architecture, Zhejiang University, Hangzhou 310058, China, and also with Zhongyuan Institute, Zhejiang University, Zhengzhou 450000, China (e-mail: jinsheng@zju.edu.cn)

Xu Min is with the Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Hong Kong, China (e-mail: min.m.xu@polyu.edu.hk)

supervision is typically implemented as the imitation of recorded AVs' trajectories. This design has two practical limitations. First, if AVs' trajectories are suboptimal, aggressive, or inconsistent human driving, pure imitation can inadvertently propagate undesirable behaviors. Second, even when trajectory-level errors are comparable, the induced interaction outcomes can differ significantly, suggesting that conventional loss values may be weakly correlated with the true quality of closed-loop planning.

To address these limitations, we propose HybridLoss, an adaptive planning-oriented loss function that strengthens the coupling between prediction and planning by introducing planner-in-the-loop supervision and interaction-aware consistency. The core idea is to substitute the recorded AVs' trajectory used in planning supervision with an adaptive motion-planning reference trajectory generated by an optimization-based module. This adaptive planner comprises parameter definition, a planning optimization module, and a scenario-dependent planning selection module. HybridLoss further augments conventional prediction and planning errors with safety potential and social force terms, encouraging consistency between predicted trajectories and planned ego behaviors in terms of interaction risk, thereby promoting stable maneuvers under dense and uncertain conditions. The main contributions are summarized as follows:

(1) **Planner-in-the-loop supervision**: We introduce an adaptive motion-planning module that generates an optimized reference trajectory and substitutes it for the ground-truth trajectory in the planning loss, thereby mitigating reliance on the strict imitation of potentially suboptimal human trajectories.

(2) **Interaction-aware coupling between prediction and planning**: We integrate safety potential and social force terms to penalize interaction-risk inconsistencies arising from prediction–planning misalignment, providing a structured feedback mechanism that better aligns prediction quality with planning outcomes.

(3) **Scenario-adaptive planning strategy selection**: We develop a selection mechanism that dynamically switches among personal, global, and hybrid planning strategies based on traffic density and interaction risk, enabling adaptive trade-offs across diverse traffic environments.

(4) **Empirical validation in interactive scenarios**: We conduct evaluations on the INTERACTION dataset across heterogeneous scenario types (roundabouts, merging, and intersections) and demonstrate consistent improvements over internal loss variants, supported by ablation studies that isolate the contributions of each component.

Following the Introduction in Section I and the Literature Review in Section II, Section III details the proposed methodology, covering key parameter definitions, the adaptive motion-planning approach, and the formulation of HybridLoss. Section IV presents a comprehensive Experimental Analysis, and Section V concludes the study.

II. Literature Review

A. Prediction Loss

Learning-based prediction relied on deterministic regression losses, which represent a straightforward approach but tend to average over multiple possible trajectories. To address multimodality, subsequent research introduced probabilistic and generative objectives, encompassing adversarial and latent-variable formulations. For instance, Social GAN uses adversarial training coupled with a diversity-oriented objective to generate socially acceptable multi-modal trajectories [3]. In the context of AV and heterogeneous-agent settings, Trajectron++ extends conditional generative prediction by incorporating dynamic feasibility and optional conditioning on AVs' planning, underscoring the importance of structured probabilistic modeling for downstream decision-making [4].

More recent works increasingly integrate map structure and interaction modeling through graph-based objectives and representations. LaneGCN models agent–map and agent–agent interactions using lane-graph convolutions, typically optimized by multi-hypothesis regression objectives over candidate trajectories [5]. VectorNet substitutes rasterized map inputs with vectorized polylines and hierarchical graph aggregation, coupled with negative log-likelihood to enhance representational efficiency and accuracy [6]. Complementary to these representational advances, one method frames motion forecasting as classification over trajectory sets, enabling stable multi-modal training via cross-entropy over anchors supplemented by residual regression, such as MultiPath [7] and CoverNet [8]. Target-driven formulations further decompose prediction into goal prediction followed by target-conditioned trajectory generation, as exemplified by TNT [9] and its extensions such as DenseTNT [10]. Transformer-based approaches advance this paradigm by leveraging global context and attention mechanisms for multi-agent motion, e.g., MTR (query-based multimodal prediction and refinement) [11] and Wayformer (efficient attention for large-scale motion forecasting) [12].

However, the dominant training signal in prediction remains grounded in data imitation: losses primarily reward proximity to logged trajectories or likelihood under a learned distribution. Consequently, forecasting objectives can be effectively minimized while still producing behaviors that are misaligned with downstream planning requirements in closed-loop operation, especially in mixed traffic where interaction outcomes outweigh pointwise similarity. Therefore, this research cannot reliably supervise or evaluate AVs' behavior when the optimal maneuver deviates from the recorded human trajectory in competitive mixed-traffic scenarios.

B. Planning Loss

Learning-based planning mainly employs Imitation Learning (IL) that regresses planned waypoints from demonstrations. Classical E2E driving approaches utilize supervised losses over waypoints. Conditional Imitation Learning (CIL) extends this paradigm by conditioning on high-level commands to resolve multimodal intent at intersections [13]. To mitigate distribution shift, dataset aggregation methods such as DAgger provide a principled framework for iteratively collecting corrective supervision [14]. In more scalable AV pipelines, planners predict a future

AVs' trajectory and are trained by planning losses, occasionally incorporating multi-hypothesis variants.

Industry-scale learning-based planners further consider imitation objectives with constraints and penalties. ChauffeurNet leverages demonstrations while incorporating losses related to traffic rules and safety indicators, effectively shaping the learned policy beyond pure regression [15]. Imitative Models bridge imitation and goal-directed planning by learning a probabilistic model of desirable behavior and optimizing trajectories under flexible goal objectives [16]. Other methods enhance planning robustness by employing supervision and teacher–student transfer. For example, Learning by Cheating trains a privileged expert-like agent and then distills this knowledge into a vision-based policy [17]. More recently, planners such as PlanT utilize object-level representations and imitation objectives to improve interpretability and efficiency, while remaining fundamentally demonstration-driven [18].

Across these methodologies, safety and interaction are often incorporated through collision, lane departure, comfort. For modeling social interactions, classical formulations like the Social Force Model provide a general mechanism to formulate proximity and interaction into potential-based penalties, which can be adapted as regularizes for socially-aware motion [19][20][21]. However, these terms typically function as auxiliary components with heuristically tuned weights, and they do not directly resolve the core ambiguity: the recorded human plan is not necessarily the optimal plan for the AV's objectives and constraints. Consequently, planning loss may (i) propagate suboptimal or aggressive human behaviors and (ii) fail to provide an adaptive "planning-consistent" supervisory target.

C. Differentiable E2E Prediction-Planning Integration Methods

A trend toward closing the prediction–planning loop involves embedding planners into learning pipelines so that gradients directly reflect downstream planning quality. On the optimization aspect, differentiable optimization layers such as OptNet introduce differentiation through constrained quadratic programs, enabling neural networks to learn parameters under constraints [22]. Moreover, differentiable Model Predictive Control (MPC) formulates MPC as a differentiable policy class and backpropagates through the Karush–Kuhn–Tucker (KKT) conditions of the controller's fixed point, allowing E2E of costs and dynamics within a control loop [23]. These methods motivate differentiable planners as a mechanism to couple learning with constraint satisfaction and stability.

In AVs, differentiable and integrated predictor–planner frameworks apply these principles to trajectory generation and interaction. Differentiable Integrated Prediction and Planning (DIPP) unifies prediction and planning through a differentiable nonlinear optimizer, enabling the joint training of prediction and planning within a single computational graph [24]. Related interactive neural planning frameworks explicitly model multi-agent coupling in Bird Eye-View (BEV) [25]. Neural Motion Planner (NMP) approaches also emphasize interpretable cost-based selection among candidate trajectories while being trained E2E [26]. Moreover, unified frameworks orchestrate perception, prediction, and planning tasks hierarchically with planning-oriented coordination, improving cross-module consistency [27]. Similarly, ST-P3 jointly learns spatial–temporal representations for perception, prediction, and planning in a unified pipeline, typically supervised via imitation-style planning targets [28].

Overall, although differentiable E2E methods improve coupling, the training objective typically comprises (i) an imitation-style planning loss to match an expert ego trajectory, supplemented by (ii) weighted auxiliary penalties for safety, comfort, and interaction. In dense mixed traffic, where optimal AVs' behavior may intentionally deviate from the logged human action to achieve stable negotiation, such objectives can remain conceptually misaligned: gradients encourage reproducing demonstrations rather than learning a planner-consistent optimum under the model's own predicted world.

D. Summary

Overall, prior work has made substantial progress in (i) multimodal, interaction-aware motion forecasting losses, (ii) imitation-driven planning losses with safety regularization, and (iii) differentiable integrated predictor–planner pipelines. However, AVs' planning supervision is typically anchored to the recorded human trajectory, with auxiliary safety and interaction losses functioning primarily as post-hoc regularizers. This creates a fundamental misalignment in mixed traffic, where "optimal" AV behavior may systematically diverge from the demonstration. Existing methodologies typically define prediction and planning objectives to coerce models to approximate ground truth data as closely as possible. However, this imitation-centric approach overlooks two critical limitations. First, minimizing standard losses does not guarantee the acquisition of optimal driving policies. Second, traditional loss magnitudes may not accurately serve as reliable proxies for the quality of iterative training outcomes. This paper aims to address two questions:

(1) **Do traditional loss functions ensure the learning of optimal driving behaviors?** In current research, loss functions incorporate prediction and planning components, with the latter focused on minimizing the displacement between AV-planned trajectories and real driver data. However, if ground truth originates from an aggressive or inexperienced driver, the model may inadvertently propagate unexpected and potentially unsafe maneuvers, thereby elevating risks to individual and collective traffic safety.

(2) **Do traditional loss definitions accurately evaluate training quality?** Current objectives often assume that lower prediction and planning loss values strictly correlate with superior performance. However, trajectories yielding equivalent loss values may exhibit spatial deviations in opposing directions relative to the ground truth, leading to divergent impacts on decision-making despite similar metric scores. As a result, standard prediction and planning losses do not always correlate monotonically with closed-loop system performance.

To resolve these challenges, this paper proposes HybridLoss, which substitutes the recorded trajectory in the planning loss with a target generated by an adaptive motion-planning module. By reformulating the evaluation criteria for prediction, planning, and overall performance, HybridLoss effectively mitigates these two scientific issues, advancing both AV model adaptability and the robustness of learned driving behaviors.

III. Methodology

This section is organized into three parts: definitions of key parameters, the adaptive motion-planning method addressing Question 1, and the loss function based on the adaptive motion-planning method, which addresses Question 2.

A. Definition of Key Parameters

This paper considers two types of agents: AVs and HVs. AVs receive information from the perception module, particularly the trajectory data of nearby HVs with which they interact. HVs, on the other hand, are controlled by human drivers. To support the analysis of the research questions, key parameters are defined in Table I.

TABLE I
DEFINITION OF PARAMETERS

Parameter	Description
AVs	Autonomous Vehicles
HVs	Human Vehicles
x_a, y_a	Lateral and longitudinal positions of AVs (m, m)
x_h, y_h	Lateral and longitudinal positions of HVs (m, m)
v_a, v_h	Speed of AVs and HVs (m/s)
a_a, a_h	Acceleration of AVs and HVs (m/s^2)
Δt	Updated frequency of trajectory (s)
T	Planning duration of trajectory (s)
$\boldsymbol{s}_a^{\mathbf{plan}}(t)$	Planning trajectory of AVs at t
$\boldsymbol{s}_a^{\mathbf{real}}(t)$	Ground-truth trajectory of AVs at t
$\boldsymbol{s}_h^{\mathbf{pred}}(t)$	Predicting trajectory of HVs at t
$\boldsymbol{s}_h^{\mathbf{real}}(t)$	Ground-truth trajectory of HVs at t
$\boldsymbol{S}_a^{\mathbf{plan}}(t)$	Planning trajectory of AVs from 0 to T
$\boldsymbol{S}_a^{\mathbf{real}}(t)$	Ground-truth trajectory of AVs from 0 to T
$\boldsymbol{S}_h^{\mathbf{pred}}(t)$	Predicting trajectory of HVs from 0 to T
$\boldsymbol{S}_h^{\mathbf{real}}(t)$	Ground-truth trajectory of HVs from 0 to T
$sf_{h\to a}$	Social force value from HVs to AVs (Interaction)
$sp_{h\to a}$	Safety potential value from HVs to AVs (Safety)
TD	Traffic density of all interaction agents
SD	Speed deviation of all interaction agents

B. Adaptive Motion-Planning Method

In previous research, loss functions typically incorporate both prediction and planning components, with the planning loss focused on aligning the AV's planned trajectories with real driver trajectory data. However, when real trajectory data is derived from an aggressive or inexperienced driver, the model may unintentionally learn suboptimal or unsafe planning decisions, thereby increasing the risks to both individual and collective traffic safety. As a result, when E2E models rely on the traditional planning loss function, the neural network may learn both optimal and suboptimal driving behaviors. This can lead to AVs generating unsatisfactory trajectories in certain scenarios due to uncertainties inherent in the learning process. To address this issue, this section introduces an adaptive motion-planning trajectory to replace the real AV trajectory in the planning loss function, enabling AVs to learn effective decision-making and motion planning more efficiently and accurately. The adaptive motion-planning method comprises two key components: the planning optimization module and the planning selection module. (Fig. 1).

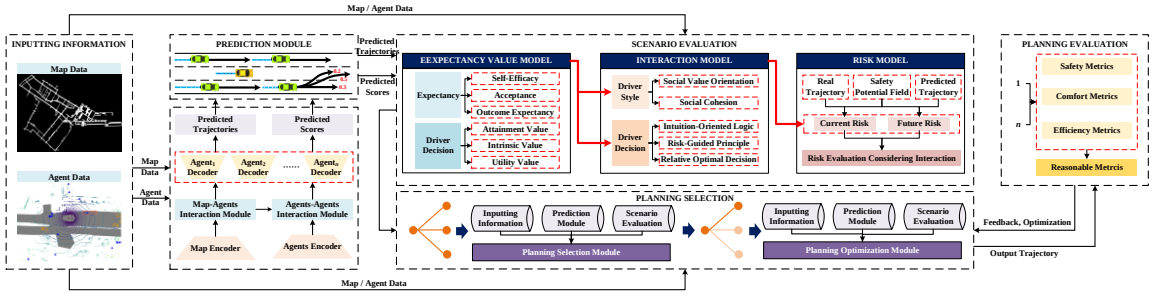

Fig. 1. Framework of the adaptive motion-planning method

1) Planning Optimization Module

In the early stages of AV development, it was relatively straightforward to define an objective function that combined three classic loss functions: safety, comfort, and efficiency. However, two critical questions arise: First, what should the true objectives for AVs be in different driving scenarios? Second, is it reasonable for AVs to simply mimic real trajectory data? To address the first question, early AV technologies primarily focused on ensuring safe and efficient driving on public roads. Regarding the second question, while a rule-based approach allows AVs to efficiently learn and replicate desired trajectories, there is a risk that AVs may inadvertently adopt aggressive driving behaviors from real-world trajectory data. Building on this analysis, this section introduces three planning strategies for E2E learning across various traffic scenarios: personal planning, global planning, and hybrid planning. The personal planning strategy is suited for low-interaction, low-risk situations, while the global planning strategy is designed for high-interaction or high-risk environments involving multiple agents. The hybrid planning strategy is employed in scenarios characterized by moderate interaction or risk.

1) Personal Planning Strategy. It is assumed that the AV generates multiple candidate trajectories, from which the optimal trajectory is selected based on the lowest cost value. The cost function for the AV's planned trajectory integrates loss functions for safety, comfort, and efficiency. The fundamental structure of these loss functions is outlined below.

Safety loss function c_s. c_s is constructed by the safety potential value and social force value from HV to AV. The equation is shown in (1):

$$c_s = \frac{T}{N \cdot \Delta t} \sum_{t=0}^{T} \sum_{i=1}^{N} \frac{sf_{h_i \to a}^{t}}{\sum_{i=1}^{N} sf_{h_i \to a}^{t}} \frac{sp_{h_i \to a}^{t}}{\max \boldsymbol{sp}_{\boldsymbol{h_i \to a}}} \tag{1}$$

where, T is the duration time for planning trajectory, in s, Δt is the update frequency for planning trajectory, in s, N is the number of HV, $sf_{h_i \to a}^{t}$ is the social force value from HV_i to AV at timestamp t, $sp_{h_i \to a}^{t}$ is the safety potential value from HV_i to AV at timestamp t, $\boldsymbol{sp}_{\boldsymbol{h_i \to a}}$ is the list of safety potential values from HV_i to AV during the range from 0 to T.

Comfort loss function c_c. c_c is constructed by the comprehensive analysis of speed v_a (m/s), acceleration a_a (m/s^2), and jerk j_a (m/s^3) of AV's candidate trajectory. The equation is shown in (2):

$$c_c = w_v \int_0^T v_a^2 dt + w_a \int_0^T a_a^2 dt + w_j \int_0^T j_a^2 dt \tag{2}$$

where w_v, w_a, and w_j are the weight values of v_a, a_a, and j_a, respectively, $w_v + w_a + w_j = 1$.

Efficiency loss function c_e. c_e is constructed by the comparison between the real speed and expected speed [29] of AV. The equation is shown in (3):

$$c_e = \frac{T}{\Delta t} \sum_{t=0}^{T} \text{sig.}\left(v_a^t - v_{a,\text{exp}}\right) - 1 \tag{3}$$

where $v_{a,\text{exp}}$ is the expected speed, in m/s and v_a^t is the real speed of AV at timestamp t.

Therefore, the equation of cost function is shown in (4):

$$C_{\text{personal}} = w_s \cdot c_s + w_c \cdot c_c + w_e \cdot c_e \tag{4}$$

where, C_{personal} is the total cost value of candidate trajectory for AV, w_s, w_c, and w_e are the weights of c_s, c_c, and c_e, respectively, $w_s + w_c + w_e = 1$. w_s, w_c, w_e are modified by [30].

2) Global Planning Strategy. The personal planning strategy focuses on the AV's own needs, making it suitable for low-risk scenarios. However, in high-density traffic, where the risk of accidents is elevated due to interactions between the AV and HVs, a global planning strategy is necessary. This strategy analyzes the distribution of traffic risk across all agents and aims to minimize overall risk, ensuring safe interactions in complex environments. Safety is prioritized in this approach because maintaining safe distances and minimizing collisions takes precedence over efficiency or comfort in high-risk traffic scenarios. Although this may result in some vehicles experiencing higher risk than others, the primary goal is to ensure collective safety for all agents. The equation is shown in (5)

$$C_{\text{global}} = \sigma\left(c_{s_a}, \boldsymbol{c}_{\boldsymbol{s_h}}^{*}\right) \tag{5}$$

where, c_{s_a} is the safety loss value for AV. $\boldsymbol{c}_{\boldsymbol{s_h}}^{*}$ is the safety loss value set for all HVs which have interactions with AV. $\sigma(c_{s_a}, \boldsymbol{c}_{\boldsymbol{s_h}}^{*})$ is the standard deviation of c_{s_a} and $\boldsymbol{c}_{\boldsymbol{s_h}}^{*}$. C_{global} is the cost value for global planning strategy.

3) Hybrid Planning Strategy. As for hybrid planning strategy, C_{personal} and C_{global} are both considered to meet part personal requirements and part global requirements. The hybrid planning strategy could be constructed by two weight factors, w_{personal} and w_{global}. The function is shown in (6):

$$C_{\text{hybrid}} = w_{\text{personal}} \cdot C_{\text{personal}} + w_{\text{global}} \cdot C_{\text{global}} \tag{6}$$

2) Planning Selection Module

Three types of motion planning strategies are defined for selection: personal motion-planning, global motion-planning, and hybrid motion-planning. Personal motion-planning generates a trajectory that prioritizes the AV's own requirements. Global motion-planning focuses on generating a trajectory that minimizes overall safety risks by considering the collective needs of all agents. Hybrid motion-planning combines both the AV's individual requirements and the broader needs of other agents. To assist AVs in selecting the most appropriate motion-planning strategy for different scenarios, a selection metric is proposed, which incorporates both global and personal parameters.

First, *TD* and *SD* are critical factors in selecting an appropriate motion-planning strategy. As traffic density increases, interactions among agents become more pronounced, creating conditions that favor global optimization. In such high-density scenarios, it becomes more challenging to prioritize the AV's individual requirements, making global planning the preferred approach. Additionally, *SD*, which quantifies the variation in speed among interacting agents, plays a crucial role in influencing the dynamics of vehicle interactions in densely packed environments. To capture the combined effects of *TD* and *SD*, we introduce a new metric, defined by (7):

$$\alpha_{gp} = w_{td} \cdot \frac{TD}{TD_{\max}} + w_{sd} \cdot \frac{SD}{SD_{\max}} \tag{7}$$

where, w_{td} is the weight for traffic density, w_{sd} is the weight for speed deviation, $TD_{\max}$ is the maximum value of traffic density [31], $SD_{\max}$ is maximum value of speed deviation [32].

Secondly, in the personal parameters, the lateral position, longitudinal position, speed, acceleration, social force value, and social potential value are calculated for HV and AV. We need to evaluate the distribution of *sf* and *sp* for each vehicle. *sf* represents social interaction so that the larger value of s_f means the more strength interaction. The same phenomenon could be also applied in *sp* for each vehicle. Therefore, we propose a new metric related to *sf* and *sp*. As for AV, $sf_{h \to a}$ represents the social force value from HV to AV, $sp_{h \to a}$ represents the safety potential value from HV to AV. Under the same $sf_{h \to a}$, the increasing of $sp_{h \to a}$ will lead to a higher risk. Under the same $sp_{h \to a}$, the increasing of $sf_{h \to a}$ will lead to a higher risk. The equation is shown in (8):

$$\alpha_a = \sum_{i=1}^{N} \frac{sf_{h_i \to a}}{\sum_{i=1}^{N} sf_{h_i \to a}} \cdot sp_{h_i \to a} \tag{8}$$

where, α_a is the metric considering s_p and s_f for AV, $sf_{h_i \to a}$ is the social force value from HV_i to AV, $sp_{h_i \to a}$ is the safety potential value from HV_i to AV.

Additionally, the impact of interactions among different agents is also relevant to HVs. Therefore, we can develop a specific metric for each HV to capture the dynamics of its interactions with surrounding agents. This metric will help quantify the HV's role in traffic flow and its influence on the overall driving environment. The equation is shown in (9):

$$\alpha_{h_i} = p_1 \cdot \sum_{j \neq i}^{N} \frac{sf_{h_j \to h_i}}{\sum_{j \neq i}^{N} sf_{h_j \to h_i}} \cdot sp_{h_j \to h_i} + p_2 \cdot sp_{h_j \to h_i} \tag{9}$$

where, p_1 and p_2 is the weights, $sf_{h_j \to h_i}$ is the social force value from HV_j to HV_i, $sp_{h_j \to h_i}$ is the safety potential value from HV_j to HV_i.

The logic to select planning strategy is shown in Fig. 2. First of all, if $\alpha_{gp} > \alpha_{gp,\max}$, global planning trajectory will be selected as planning strategy without considering α_a and α_{h_i}. Otherwise, we will make comparisons between α_a and $\alpha_{a,\max}$, α_{h_i} and $\alpha_{h_i,\max}$. Whatever $\alpha_a > \alpha_{a,\max}$ or $\alpha_{h_i} > \alpha_{h_i,\max}$, we will select global planning trajectory while the next judgement will be calculated. If $\sigma(\alpha_a, \boldsymbol{\alpha_h}) > \sigma_{\max}$, hybrid planning trajectory will be selected as the strategy otherwise the personal planning trajectory will be selected. When we select hybrid planning, w_{personal} and w_{global} could be updated by (10)-(13):

$$w'_{\mathrm{personal}} = \left(1 + \frac{\alpha_a}{\alpha_{a,\max}}\right)^{\sigma(\alpha_a, \boldsymbol{\alpha_h})} w_{\mathrm{personal}} \tag{10}$$

$$w'_{\mathrm{global}} = \left(1 + \frac{\alpha_{gp}}{\alpha_{gp,\max}}\right) w_{\mathrm{global}} \tag{11}$$

$$w^{\mathrm{update}}_{\mathrm{personal}} = \frac{w'_{\mathrm{personal}}}{w'_{\mathrm{personal}} + w'_{\mathrm{global}}} \tag{12}$$

$$w^{\mathrm{update}}_{\mathrm{global}} = \frac{w'_{\mathrm{global}}}{w'_{\mathrm{personal}} + w'_{\mathrm{global}}} \tag{13}$$

where, w'_{personal} and w'_{global} are the calculated weights in the hybrid planning, $w^{\mathrm{update}}_{\mathrm{personal}}$ and $w^{\mathrm{update}}_{\mathrm{global}}$ are the final weights in the hybrid planning. $\alpha_{a,\max}$ is the maximum of α_a, $\alpha_{h_i,\max}$ is the maximum of α_{h_i}, $\alpha_{gp,\max}$ is the maximum of α_{gp}, $\sigma(\alpha_a, \boldsymbol{\alpha_h})$ is the standard deviation among AV and HVs which have interactions with AV. $\boldsymbol{\alpha_h} = \{\alpha_{h_1}, \alpha_{h_2}, \ldots, \alpha_{h_N}\}$.

Algorithm 1 Planning Selection Module

Input: Parameters $\alpha_a, \alpha_{h_i}, \alpha_{gp}, \alpha_{a,\max}, \alpha_{h_i,\max}, \alpha_{gp,\max}, \sigma_{\max}$
Output: Selected planning strategy and, if hybrid plan is selected, updated weights $w^{\mathrm{update}}_{\mathrm{personal}}, w^{\mathrm{update}}_{\mathrm{global}}$

Initialize: $w_{\mathrm{personal}}, w_{\mathrm{global}}$
Step 1: Check if $\alpha_{gp} > \alpha_{gp,\max}$
if $\alpha_{gp} > \alpha_{gp,\max}$ **then**
  Select global planning trajectory as the strategy
  Return selected strategy
else
  Proceed to Step 2
end if
Step 2: Compare α_a with $\alpha_{a,\max}$ and α_{hi} with $\alpha_{hi,\max}$
if $\alpha_a > \alpha_{a,\max}$ **or** $\alpha_{hi} > \alpha_{hi,\max}$ **then**
  Select global planning trajectory and calculate the next judgment
  Return selected strategy
else
  Proceed to Step 3
end if
Step 3: Calculate $\sigma(\alpha_a, \alpha_h)$ and compare with $\sigma_{\max}$
if $\sigma(\alpha_a, \alpha_h) > \sigma_{\max}$ **then**
  Select hybrid planning trajectory as the strategy
  Step 4: Update w_{personal} and w_{global}:
  $w'_{\mathrm{personal}} \leftarrow \left(1 + \frac{\alpha_a}{\alpha_{a,\max}}\right) \sigma(\alpha_a, \alpha_h) w_{\mathrm{personal}}$
  $w'_{\mathrm{global}} \leftarrow \left(1 + \frac{\alpha_{gp}}{\alpha_{gp,\max}}\right) w_{\mathrm{global}}$
  Step 5: Calculate updated weights:
  $w^{\mathrm{update}}_{\mathrm{personal}} \leftarrow \frac{w'_{\mathrm{personal}}}{w'_{\mathrm{personal}} + w'_{\mathrm{global}}}$
  $w^{\mathrm{update}}_{\mathrm{global}} \leftarrow \frac{w'_{\mathrm{global}}}{w'_{\mathrm{personal}} + w'_{\mathrm{global}}}$
  Return hybrid strategy and updated weights
else
  Select personal planning trajectory as the strategy
  Return selected strategy
end if

Fig. 2. Pseudocode of the planning selection module

3) Relation to Desensitized Trajectory Optimization

Desensitized-aware trajectory optimization [33] improves planning stability by introducing robustness constraints within the optimization problems, which is solved for each planning instance. Adaptive motion-planning component is related in its goal but differs fundamentally in its role and integration: A scenario-adaptive planner is employed to generate a reference trajectory s_a^{adap} that replaces the ground-truth trajectory s_a^{real} in the planning loss during training. This "planner-in-the-loop loss shaping" mitigates imitation of suboptimal human demonstrations and steers E2E learning toward rational planning behaviors grounded in safety, comfort, and efficiency. Moreover, unlike robust optimization that focuses on the AV's trajectory alone, HybridLoss introduces interaction-aware penalties (safety potential and social force discrepancies) that explicitly couple prediction errors with planning risk, enforcing a closed-loop consistency between prediction and planning under mixed-traffic interactions.

C. Loss Function Based on Adaptive Motion-Planning

Traditionally, the loss function comprises two main components: the prediction and the planning loss function. The prediction loss function calculates the deviation error between the predicted trajectories and real trajectories of HVs, while the planning loss function measures the deviation between the AV's planning and real trajectories. However, there is a weak correlation between the prediction and planning loss functions, making it difficult to determine whether a lower loss value reliably indicates better performance. Additionally, integrating the planning loss function with the adaptive motion-planning method discussed earlier remains challenging. To address these issues, this section introduces a new loss function, HybridLoss, which replaces real trajectory data in the planning loss function with an adaptive motion-planning module.

The equation of prediction loss function is shown in (14):

$$L_{h_i} = \sqrt{\frac{T}{\Delta t} \sum_{t=0}^{T} \frac{\left\| \mathbf{s}_{h_i}^{\mathbf{pred}}(t) - \mathbf{s}_{h_i}^{\mathbf{real}}(t) \right\|^2}{2}} \tag{14}$$

where, L_{h_i} is the root mean square error between prediction trajectory and real trajectory of HV$_i$, $\boldsymbol{s}_{\boldsymbol{h_i}}^{\mathbf{pred}}(\boldsymbol{t})$ and $\boldsymbol{s}_{\boldsymbol{h_i}}^{\mathbf{real}}(\boldsymbol{t})$ are the prediction and real trajectories of HV$_i$ at *t*.

The equation of planning loss function is shown in (15):

$$L_a = \sqrt{\frac{T}{\Delta t}\sum_{t=0}^{T}\frac{\left\|\boldsymbol{s}_{\boldsymbol{a}}^{\mathbf{pred}}(\boldsymbol{t}) - \boldsymbol{s}_{\boldsymbol{a}}^{\mathbf{real}}(\boldsymbol{t})\right\|^2}{2}} \tag{15}$$

where, L_a is the root mean square error between planning trajectory and real trajectory of AV, $\boldsymbol{s}_{\boldsymbol{a}}^{\mathbf{pred}}(\boldsymbol{t})$ and $\boldsymbol{s}_{\boldsymbol{a}}^{\mathbf{real}}(\boldsymbol{t})$ are the prediction and real trajectory of AV at *t*.

To improve the performance of planning loss function, we replace real trajectory data of AV in the planning loss function with an adaptive motion-planning module. The updated formulation of La is shown in (16):

$$L_a^{'} = \sqrt{\frac{T}{\Delta t}\sum_{t=0}^{T}\frac{\left\|\boldsymbol{s}_{\boldsymbol{a}}^{\mathbf{pred}}(\boldsymbol{t}) - \boldsymbol{s}_{\boldsymbol{a}}^{\mathbf{adap}}(\boldsymbol{t})\right\|^2}{2}} \tag{16}$$

where, $L_a^{'}$ is the updated planning loss function, $\boldsymbol{s}_{\boldsymbol{a}}^{\mathbf{adap}}(\boldsymbol{t})$ is the planning trajectory of AV from adaptive motion-planning module at *t*.

The final loss function is shown in (17).

$$L_{\text{initial}} = w_a \cdot L_a^{'} + \sum_{i=1} w_{h_i} \cdot L_{h_i} \tag{17}$$

where w_a and w_{h_i} are the weights of planning and prediction loss function.

Prediction loss function plays a critical role by ensuring an accurate understanding of the behaviors of HVs. Without precise predictions of HVs' trajectories, the planning module would lack essential information about the future states of other road users, potentially leading to suboptimal or unsafe planning decisions. This deficiency could result in the AV failing to account for the actions of other agents on the road. The planning loss function evaluates how closely the AV's planned trajectory aligns with the ideal, ground-truth trajectory, thereby optimizing the vehicle's path in terms of safety, comfort, and efficiency. While the prediction and planning loss functions are individually essential for evaluating their respective modules, an important question arises: can E2E performance be accurately assessed through a simple combination of prediction and planning loss values?

Example 1: As shown in Fig. 3, there are two vehicles: an AV and a HV. The AV's actual trajectory points over three timestamps are [(0, 0), (0.5, 0.3), (0.4, 0.7)], while the HV's actual trajectory points are [(2.8, 3.0), (3.0, 4.5), (3.3, 4.3)]. The AV has two candidate planned trajectories. The first planned trajectory is [(0, 0), (0.7, 0.5), (0.6, 0.9)], and the second planned trajectory is [(0, 0), (0.3, 0.1), (0.2, 0.5)]. The values of L_a remain the same under different planning outcomes. Similarly, there are two predicted trajectories for the HV. The first predicted trajectory is [(2.8, 3.0), (3.3, 4.8), (3.5, 4.5)], and the second predicted trajectory is [(2.8, 3.0), (2.7, 4.2), (3.1, 4.1)]. The values of L_h also remain the same under different prediction outcomes. In a traditional loss function, these two scenarios would yield the same performance evaluation. However, the first scenario presents a higher risk than the second. Therefore, it is necessary to explore the relationship between prediction loss and planning loss to address such discrepancies.

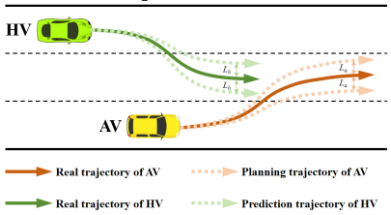

Fig. 3. The prediction and planning diagram of AV

First of all, the lower values of prediction loss and planning loss lead to a better performance. Moreover, in order to output the planning trajectory for AV, the weight of planning loss function needs to be higher than prediction loss function $w_a > \sum_{i=1}^{N} w_{h_i}$.

Secondly, we could calculate the prediction and real social force value at each future timestamp ($sf_{h_i\to a}^{\text{pred}}(t)$, $sf_{h_i\to a}^{\text{real}}(t)$), and the prediction and real safety potential value at each future timestamp ($sp_{h_i\to a}^{\text{pred}}(t)$, $sp_{h_i\to a}^{\text{true}}(t)$). $L_a^{'}$ and L_{h_i} could evaluate the performances of prediction module and planning module separately, which is used to guarantee the output trajectories under a relatively reasonable range. Under the same $L_a^{'}$ and L_{h_i}, the lower deviation error between $sf_{h_i\to a}^{\text{pred}}(t)$ and $sf_{h_i\to a}^{\text{real}}(t)$ or between $sp_{h_i\to a}^{\text{pred}}(t)$ and $sp_{h_i\to a}^{\text{real}}(t)$ will lead to a better performance of the combination of predicted trajectory of HV and the planned trajectory of AV. Overall, $L_a^{'}$ and L_{h_i} are used to make sure the robustness of training results, and $sf_{h_i\to a}^{\text{pred}}(t)$, $sf_{h_i\to a}^{\text{real}}(t)$, $sp_{h_i\to a}^{\text{pred}}(t)$, $sp_{h_i\to a}^{\text{real}}(t)$ are used to guide the E2E model to find the right direction to evaluate the performance of training results. Therefore, the basic loss function which considers the relationship between the prediction loss function and planning loss function is shown in (18)(19):

$$L_{sf_i} = \sqrt{\frac{T}{\Delta t}\sum_{t=0}^{T}\frac{\left\|sf_{h_i\to a}^{\text{pred}}(t) - sf_{h_i\to a}^{\text{real}}(t)\right\|^2}{2}} \tag{18}$$

$$L_{sp_i} = \sqrt{\frac{T}{\Delta t}\sum_{t=0}^{T}\frac{\left\|sp_{h_i\to a}^{\text{pred}}(t) - sp_{h_i\to a}^{\text{real}}(t)\right\|^2}{2}} \tag{19}$$

HybridLoss between AV and HVs is shown in (20):

$$L_{\text{hybrid}} = \sum_{i=1}^{N}\sum_{t=0}^{T}\frac{\alpha_{sp_i} L_{sp_i}(t)}{e^{\lambda t}} \cdot L_{\text{initial}}^{\alpha_{sf_i} L_{sf_i}(t)}, w_a > \sum_{i=1} w_{h_i} \tag{20}$$

where, L_{hybrid} is the hybrid loss function about AV and HVs, α_{sp_i} is the parameter of L_{sp_i}, λ is decay rate, the higher value of λ means the higher speed of decay, α_{sf_i} is the parameter of L_{sp_i}.

The differences between desensitized trajectory optimization and HybridLoss are shown in Table II:

TABLE II
COMPARISONS WITH TRAJECTORY OPTIMIZATION METHODS

Dimension	Desensitized trajectory optimization	HybridLoss (ours)
Goal	Stability via sensitivity penalties	Rational E2E behavior shaping
Robustness source	External planner constraints	Internal loss function design
Supervision	Standard ground truth	Planner-guided reference
Adaptivity	Static constraint tuning	Dynamic strategy selection
Coupling	Decoupled risk handling	Explicit interaction coupling

IV. EXPERIMENTAL ANALYSIS

A. Experimental Settings

1) Dataset and Experimental Protocols

(i) **Data description and split**. The INTERACTION dataset includes four categories of interactive driving scenarios: roundabouts, un-signalized intersections, signalized intersections, and merging/lane-changing. We utilize the INTERACTION dataset for planning strategy analysis for several reasons. First, its high-accuracy trajectory data enhances the reliability of experimental results. Second, it encompasses a variety of high- and low-density traffic scenarios, providing extensive data for analyzing and selecting effective planning strategies. Visualizations of each scenario are presented in Fig. 4 below.

After preprocessing, the dataset comprises 27,820 training scenes and 6,659 testing scenes. Each scene includes trajectory data for a single AV and multiple surrounding vehicles (SVs), where agent states are defined by 2D Cartesian coordinates (x, y). The sampling interval is set to Δt = 0.1 s. Historical states for both AV and SVs are encoded as flattened 10-dimensional vectors, corresponding to a history horizon of T_{hist} =5 steps. The future prediction horizon for SVs extends to T_{pred} =10 steps, while the AV planning and evaluation horizon is set to T_{plan} =4 steps. The planner generates trajectories of length 5 but truncates them to 4 steps to align with available ground-truth labels. To capture multi-agent dependencies, we construct a fully connected interaction graph among SV agents for each scene and encode features using a two-layer GAT-based interaction module. Finally, per-sample min–max normalization is applied to all past and future trajectories prior to the computation of prediction, planning, and interaction losses.

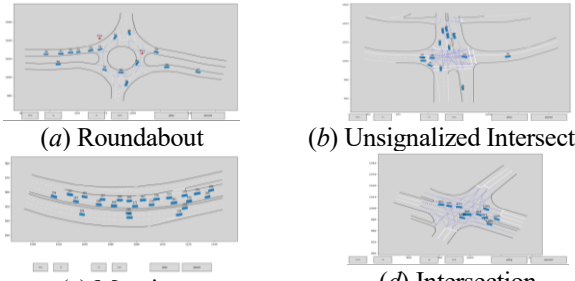

(*a*) Roundabout (*b*) Unsignalized Intersection (*c*) Merging (*d*) Intersection

Fig. 4. The visualized of different scenarios

(ii) **Training configuration**. The experiment trains for up to 15 epochs with Adam (learning rate 1e−4, weight decay 1e−5) and a StepLR schedule (step size 10, gamma 0.5). Batch size is 32. We apply gradient clipping with max norm 1.0 and use early stopping with patience 10 based on validation/training loss.

(iii) **Formalized metrics**. Let $\mathbf{p}_t^{'}$ and $\mathbf{p}_t$ be predicted and ground-truth positions at t. Metrics definition is shown below:

Average Displacement Error (ADE) is shown in (21):

$$\text{ADE} = \frac{1}{T}\sum_{t=1}^{T}\left\|\mathbf{p}_t^{'} - \mathbf{p}_t\right\|_2 \tag{21}$$

Final Displacement Error (FDE) is shown in (22):

$$\text{FDE} = \left\|\mathbf{p}_T^{'} - \mathbf{p}_T\right\|_2 \tag{22}$$

Miss Rate (MR) is a miss occurs if $\left\|\mathbf{p}_t^{'}\text{-}\mathbf{p}_t\right\|_2 > \tau_{\text{MR}}$, τ_{MR} = 2.0 m and report the percentage over all predicted points. End-to-End Prediction Accuracy (EPA) is the fraction of points with distance $\le \tau_{\text{EPA}}$ and report is a percentage, τ_{EPA} = 4.0 m. L2 error is the average AV displacement error at 1st/2nd/3rd future steps. Collision Rate (CR) is counted if AV-SV distance at a timestamp is below τ_{col}, τ_{col} = 1.0 m, and collision percentages are reported at 1st/2nd/3rd future steps.

(iv) **Hardware specifications**. The experiments are processed on a server equipped with four NVIDIA GeForce RTX 4090 GPUs (24GB VRAM each), an AMD EPYC 7542 32-Core Processor (64 threads), and 256GB DDR4 RAM.

2) Network Architecture

Our model is a three-block network (interaction, prediction, planning module). Interaction module is a graph network that encodes inter-agent relations using GATConv. Prediction module is a Transformer encoder that takes concatenated agent history embeddings and interaction features, and outputs (i) future positions via a linear head and (ii) per-timestep scores via a linear head. Planning module generates candidate trajectories using a linear candidate generator and selects minimum-cost candidate under a weighted cost (Table III).

TABLE III
NETWORK ARCHITECTURE

(*a*) Interaction module

Attribute	Details
Architecture	2-layer GNN (GATConv), ReLU, Dropout
Hyperparameters	Heads = 4, Hidden = 64, Dropout = 0.3
Input Dimension	Node feature $x \in \mathbb{R}^{N\times 10}$
Output Dimension	Interaction features $h \in \mathbb{R}^{N\times 64}$

(*b*) Prediction module

Attribute	Details
Architecture	Transformer encoder, 2 linear heads
Hyperparameters	d_{model} = 64, n_{head} = 8, layers = 3, FFN = 512, dropout = 0.1
Input Features	Past traj. embedding + interaction features
Output Dimension	Positions $\in \mathbb{R}^{T\times N\times 64}$, Score $\in \mathbb{R}^{T\times N}$

(*c*) Planning module

Attribute	Details
Architecture	Candidate generator, cost evaluation, argmin selection
Hyperparameters	$N_{\text{candidate}}$ = *20,* time steps = *5,* trainable weights
Input Features	Ego state (speed profile) + SV predicted futures
Output Dimension	Best ego plan $\in \mathbb{R}^{B\times 4\times 2}$

3) End-to-End Flow

In each scene, we construct a fully connected interaction graph among surrounding vehicles, (ii) encode interaction features with the GAT-based interaction module, (iii) predict surrounding-vehicle futures with the Transformer-based prediction module, and (iv) plan the AVs' trajectory using the planning module conditioned on predicted surrounding futures (Fig. 5).

Algorithm 1: Training with HybridLoss (Planner-in-the-loop + Interaction Consistency)

Input: Batch of scenes $\mathcal{B} = \{(S_{SV}^{past}, S_{SV}^{future}, S_{EV}^{past}, S_{EV}^{future}, \text{EV\_state})\}$
Output: Updated model parameters θ

1. Interaction Encoding:
Construct complete interaction graph $\mathcal{G}$ with `edge_index` over SV nodes.
$H \leftarrow \text{InteractionModule}(S_{SV}^{past}, \text{edge\_index})$
2. SV Prediction:
$(\hat{S}_{SV}^{future}, \hat{s}) \leftarrow \text{PredictionModule}(S_{SV}^{past}, H)$
3. EV Planning (Planner-in-the-loop):
Generate K candidate trajectories by perturbing S_{EV}^{past} with Gaussian noise $\mathcal{N}(0, \sigma)$.
for $k \leftarrow 1$ **to** K **do**
  Compute candidate cost: $C_k = w_s C_{safety} + w_c C_{comfort} + w_e C_{eff}$
end
Select best plan: $\hat{S}_{EV}^{plan} \leftarrow \arg\min_k C_k$
Truncate $\hat{S}_{EV}^{plan}$ to 4 time steps.
4. Loss Computation:
Compute base losses (SmoothL1):
$L_{pred} = \text{SmoothL1}(\hat{S}_{SV}^{future}, S_{SV}^{future})$
$L_{plan} = \text{SmoothL1}(\hat{S}_{EV}^{plan}, S_{EV}^{future})$
Interaction Measures (Safety Potential & Social Force):
Compute real metrics (sp_{real}, s_{real}) using $(S_{SV}^{future}, S_{EV}^{future})$.
Compute predicted metrics (sp_{pred}, s_{pred}) using $(\hat{S}_{SV}^{future}, \hat{S}_{EV}^{plan})$.
HybridLoss (Consistency Shaping):
$\sigma_{sp} \leftarrow \text{std}([sp_{real}, sp_{pred}])$, $\sigma_s \leftarrow \text{std}([s_{real}, s_{pred}])$
$L_{hybrid} = (\alpha_{sp}\sigma_{sp}) \cdot (1 + \lambda_{pred} L_{pred} + \lambda_{plan} L_{plan})^{(\alpha_s \sigma_s)}$
5. Optimization:
$L_{total} = L_{hybrid} + L_{pred} + L_{plan}$
Backpropagate L_{total}, apply gradient clipping.
Update θ using Adam optimizer.

Fig. 5. End-to-End flow framework

4) Planner Differentiability Clarification

Planning module uses sampling and discrete selection. Candidate trajectories and their costs are differentiable network parameters and the trainable cost weights, but the final argmin selection is non-differentiable. In our implementation, gradients propagate through the selected trajectory/cost branch, while the discrete selection itself is treated as a non-differentiable operation.

5) Hyperparameters Settings

TABLE IV
TRAINING AND EVALUATION HYPERPARAMETERS

Category	Setting
Batch size	32
Epochs	Up to 15, early stopping patience = 10
Optimizer	Adam (lr = 1e-4, weight decay = 1e-5)
LR schedule	StepLR (step size = 10, gamma = 0.5)
Gradient clipping	Max norm = 1.0
Interaction GNN	GATConv, 2 layers, heads = 4, hidden = 64, dropout = 0.3
Transformer	$d_{model} = 64$, $n_{head} = 8$, encoder layers = 3, FFN = 512, dropout = 0.1
Planner candidates	$N_{candidates} = 20$, time steps = 5
Loss weights	$\lambda_{pred} = 0.33$, $\lambda_{plan} = 0.33$, $\lambda_{hybrid} = 0.33$, $\alpha_{sp} = 1.0$, $\alpha_{sf} = 0.5$
Collision metric	Threshold = 1.0 m
MR metric	Threshold = 2.0 m
EPA metric	Threshold = 4.0 m

B. Analysis of Planning Selection Module

In this part, we will analyze the distribution of α_{gp} and $\alpha_h(\alpha_a)$.

1) Distribution of α_{gp} at Each Timestamp

The dataset includes several classic scenarios, such as roundabouts, merging, signalized intersections, and un-signalized intersections. As shown in Table V and Fig. 6, *TD*, *SD*, and α_{gp} are lowest in roundabout scenarios, whereas these parameters are highest in un-signalized intersections due to the increased interactions among agents.

TABLE V
DISTRIBUTION OF TD AND SD

Scenarios	Index	Num.	Avg.	Std.	Min.	Med.	Max.
Roundabout	*TD*	468176	0.08455	0.10994	0.00121	0.01416	0.62131
	SD		0.06329	0.06138	0.00001	0.04424	0.54507
	α_{gp}		0.14784	0.11942	0.00181	0.10660	0.87470
Merging	*TD*	2084620	0.00855	0.04026	0.00055	0.00230	0.66565
	SD		0.13408	0.11039	0.00001	0.10254	0.89204
	α_{gp}		0.14263	0.11312	0.00078	0.11062	0.94334
Signalized Intersection	*TD*	132120	0.17227	0.17465	0.00388	0.0772	0.97222
	SD		0.25782	0.19576	0.00041	0.24000	1.00000
	α_{gp}		0.43009	0.26939	0.00608	0.40171	1.45574
Unsignalized Intersection	*TD*	525854	0.11600	0.13257	0.00101	0.03221	0.92447
	SD		0.10575	0.10291	0.00001	0.07251	0.75055
	α_{gp}		0.22176	0.16147	0.00135	0.20506	1.00385

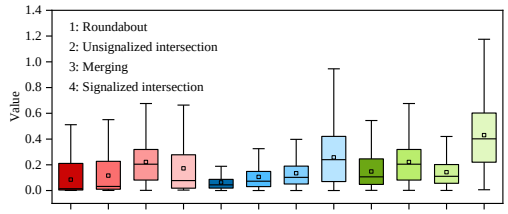

Fig. 6. The boxplot of *TD*, *SD*, α_{gp}

2) Distribution of $\alpha_h(\alpha_a)$ at Each Timestamp

As shown in Table VI and Fig. 7, the values of *sf*, *sp*, and $\alpha_h(\alpha_a)$ are lowest in roundabout scenarios, at 0.08455, 0.06329, and 0.14784, respectively, compared to other scenarios. In contrast, these values are highest in un-signalized intersections, indicating that increased interactions among agents correspond to a higher level of risk.

TABLE VI
DISTRIBUTION OF *sf*, *sp*, AND $\alpha_a(\alpha_h)$

Scenarios	Index	Num.	Avg.	Std.	Min.	Med.	Max.
Merging	*sf*	6198700	0.19679	0.22848	0.00000	0.11050	3.57310
	sp		1.81878	2.09845	0.00440	1.01400	18.6946
	α_a/α_h	2082300	2.53634	2.07564	0.11111	1.88349	17.8461
Roundabout	*sf*	581067	0.09911	0.11601	0.00000	0.06440	2.47430
	sp		1.09019	1.19978	0.00440	0.70130	10.9077
	α_a/α_h	368214	1.24155	1.22537	0.00464	0.85396	10.9077
Signalized Intersection	*sf*	230636	0.61528	0.63519	0.00000	0.43700	6.62080
	sp		4.91364	4.21328	0.00470	3.84270	20.3298
	α_a/α_h	132120	5.26953	4.06267	0.00498	4.71472	20.2784
Unsignalized Intersection	*sf*	747717	0.18171	0.26103	0.00000	0.09740	4.51380
	sp		1.62515	1.97349	0.00440	0.91280	15.4838
	α_a/α_h	534088	2.13364	1.9937	0.11112	1.49468	15.8379

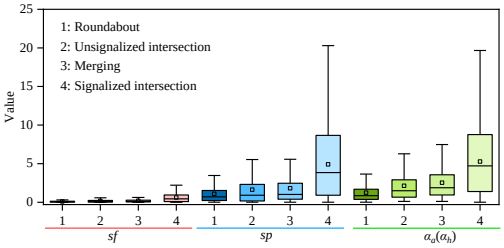

Fig. 7. The boxes of *sf*, *sp*, and α_a/α_h

C. Analysis of Adaptive Motion-Planning Method

We evaluate the performance of selection planning (our method), personal planning, global planning, hybrid planning, and human data. The personal metrics include c_s, c_c, and c_e, while the global metrics comprise the average and standard deviation of c_s, and the average value of c_e.

As shown in Table VII, the performance of c_c is highest among all methods, indicating that drivers prioritize comfort when making decisions or planning. In terms of personal metrics, selection planning outperforms in both c_s and c_e, demonstrating that our method effectively ensures both safety and efficiency. Additionally, based on global metrics, selection planning achieves a balanced distribution of requirements among agents compared to other methods. Notably, in signalized intersections, the values for c_s (avg.), c_s (*ST*), and c_e (avg.) are 2.694, 1.582, and 2.652, respectively.

TABLE VII
COMPARISON AMONG DIFFERENT METHODS

Scenario	Method	Personal Index (Pre-Norm)			Global Index (Pre-Norm)		
		c_s ↓	c_c ↓	c_e ↓	c_s (avg.) ↓	c_s (ST) ↓	c_e (Avg.) ↓
Merging	**Selection***	**0.217**	1.950	**2.914**	2.892	**1.546**	**2.538**
	Personal	0.429	5.908	5.677	3.905	12.341	5.918
	global	0.511	5.694	5.798	3.847	2.062	8.196
	Hybrid	1.020	5.534	9.159	5.774	3.084	12.779
	Human	3.535	**0.229**	5.492	**1.033**	3.247	5.628
Roundabout	**Selection***	**0.078**	1.202	**3.568**	2.163	**1.807**	2.505
	Personal	0.161	3.647	6.949	1.647	8.225	7.021
	global	0.187	3.201	7.199	2.682	2.705	7.977
	Hybrid	0.272	3.264	11.144	4.327	3.613	12.681
	Human	1.358	**0.350**	7.270	**1.618**	7.159	**2.398**
Signalized Intersection	**Selection***	**0.271**	0.719	**3.991**	**2.694**	**1.582**	**2.652**
	Personal	0.544	2.179	7.774	3.785	11.520	7.904
	global	0.588	1.936	7.911	3.593	2.108	8.577
	Hybrid	0.899	1.804	12.421	5.390	3.161	13.374
	Human	3.939	**0.438**	8.470	4.488	6.405	5.718
Unsignalized Intersection	**Selection***	**0.146**	1.320	**3.465**	2.393	**1.734**	**2.509**
	Personal	0.292	3.753	6.889	2.280	9.340	6.978
	global	0.330	3.605	6.972	3.173	2.317	8.154
	Hybrid	0.533	3.408	11.018	4.761	3.474	12.746
	Human	1.800	**0.572**	6.730	**2.092**	3.082	7.204

Moreover, we draw the box plot for each metric, which prove the best robustness in personal and global index under different classic scenarios (Figs. 8-15).

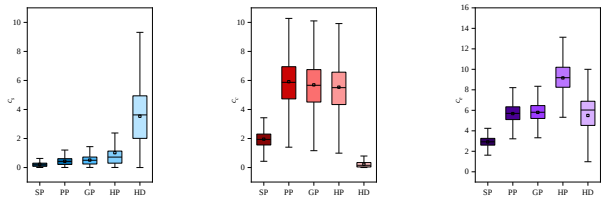

(*a*) c_s (*b*) c_c (*c*) c_e

Fig. 8. Distribution of c_s, c_c, and c_e in merging

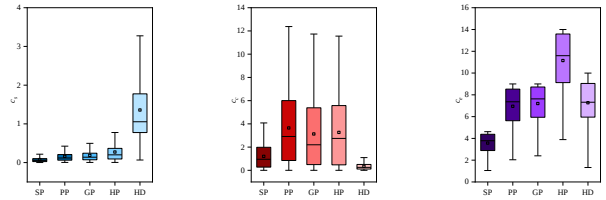

(*a*) c_s (*b*) c_c (*c*) c_e

Fig. 9. Distribution of c_s, c_c, and c_e in roundabout

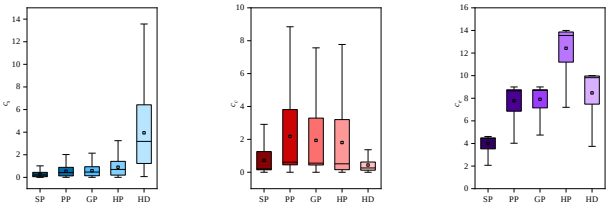

(*a*) c_s (*b*) c_c (*c*) c_e

Fig. 10. Distribution of c_s, c_c, and c_e in signalized intersection

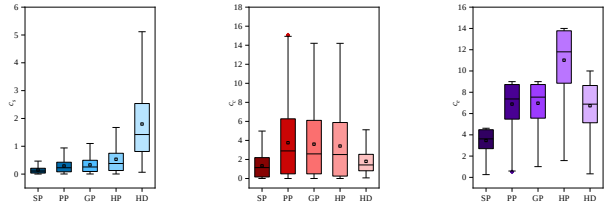

(*a*) c_s (*b*) c_c (*c*) c_e

Fig. 11. Distribution of c_s, c_c, and c_e in unsignalized intersection

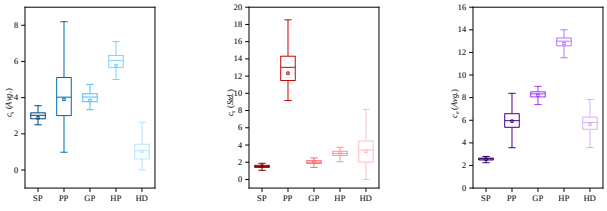

(*a*) c_s (Avg.) (*b*) c_s (ST) (*c*) c_e (Avg.)

Fig. 12. Distribution of c_s (Avg.), c_s (ST), and c_e (Avg.) in merging

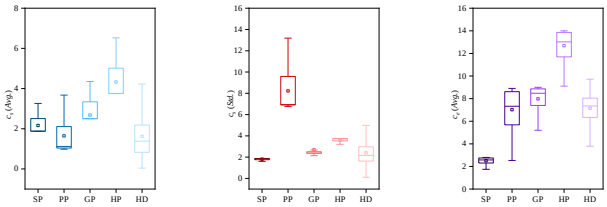

(*a*) c_s (Avg.) (*b*) c_s (ST) (*c*) c_e (Avg.)

Fig. 13. Distribution of c_s (Avg.), c_s (ST), and c_e (Avg.) in roundabout

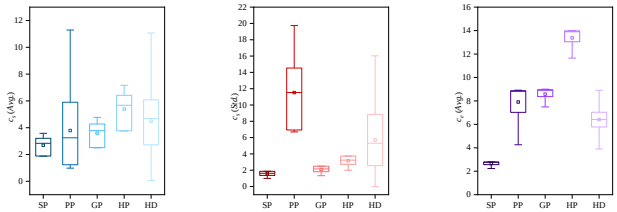

(*a*) c_s (Avg.) (*b*) c_s (ST) (*c*) c_e (Avg.)

Fig. 14 Distribution of c_s (Avg.), c_s (ST), and c_e (Avg.) in signalized intersection

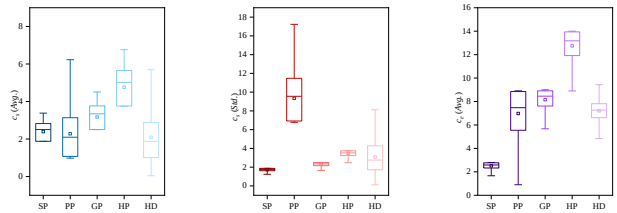

(*a*) c_s (Avg.) (*b*) c_s (ST) (*c*) c_e (Avg.)

Fig. 15. Distribution of c_s (Avg.), c_s (ST), and c_e (Avg.) in unsignalized intersection

D. Analysis of HybridLoss

In this part, we make a comparison between class loss function and HybridLoss, and ablation analysis.

1) **Comparison Analysis**

We compare four loss functions: Loss 1 (Prediction Loss and Planning Loss), Loss 2 (Prediction Loss, Planning Loss, and Cost Loss), Loss 3 (Prediction Loss, Prediction Scores, Planning Loss, and Cost Loss), and Loss 4 (Hybrid Loss, Prediction Loss, and Planning Loss). Fig. 16 shows the training loss of different E2E models under Loss 1 to Loss 4.

As shown in Table VIII, for the prediction module, the Average Displacement Error (ADE), Final Displacement Error (FDE), Miss Rate (MR), and End-to-End Prediction Accuracy (EPA) for HybridLoss are 1.11 m, 1.36 m, 0.14, and 0.28, respectively—the lowest values compared to Loss 1, Loss 2, and Loss 3. For the planning module, the averages of L2 distance and Collision Rate are 0.86 m and 0.11%, outperforming the other functions.

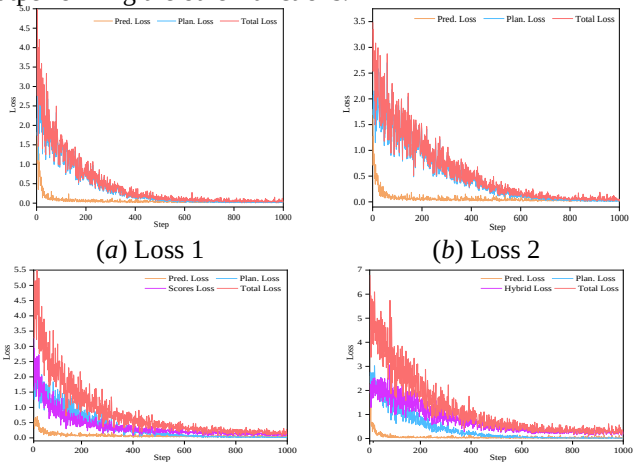

(*a*) Loss 1 (*b*) Loss 2 (*c*) Loss 3 (*d*) Loss 4

Fig. 16. Training loss

TABLE VIII
COMPARISON AMONG DIFFERENT LOSS FUNCTIONS

Loss	Prediction Module				Planning Module							
					L2↓ (m)				Col. Rate↓ (%)			
	ADE↓ (m)	FDE↓ (m)	MR↓	EPA↑	1 *s*	2 *s*	3 *s*	Avg.	1 *s*	2 *s*	3 *s*	Avg.
Loss 1	1.36±0.04	1.64±0.05	0.22±0.00	0.18±0.00	0.55±0.08	1.04±0.14	1.59±0.25	1.06±0.16	0.11±0.03	0.21±0.03	0.23±0.04	0.19±0.03
Loss 2	1.37±0.02	1.65±0.03	0.22±0.00	0.19±0.00	**0.32±0.10**	**0.81±0.31**	1.47±0.38	0.87±0.26	0.17±0.02	0.25±0.02	0.25±0.05	0.22±0.04
Loss 3	1.36±0.01	1.63±0.01	0.22±0.00	0.19±0.00	0.58±0.13	1.17±0.18	1.59±0.42	1.11±0.24	0.11±0.03	0.20±0.02	0.24±0.05	0.18±0.04
HybridLoss*	**1.11±0.03**	**1.36±0.03**	**0.14±0.00**	**0.28±0.00**	0.41±0.01	0.85±0.02	**1.33±0.04**	**0.86±0.03**	**0.07±0.01**	**0.13±0.03**	**0.14±0.04**	**0.11±0.03**

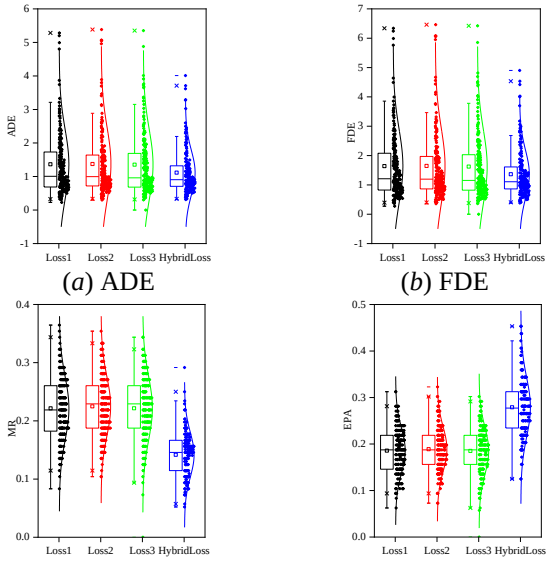

(*a*) ADE (*b*) FDE (*c*) MR (*d*) EPA

Fig. 17. The performance of prediction module

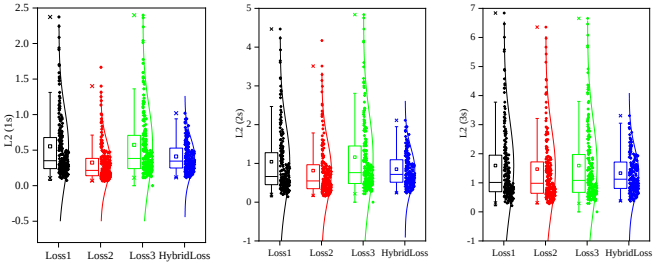

(*a*) 1 s (*b*) 2 s (*c*) 3 s

Fig. 18. Performance of planning module (1)

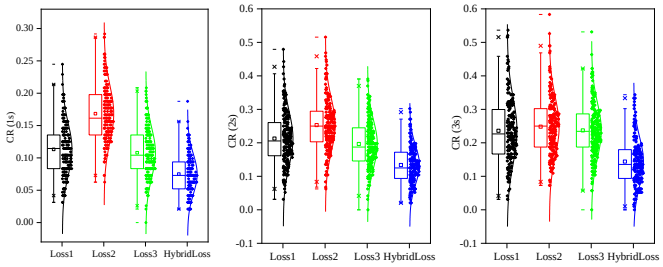

(*a*) 1 s (*b*) 2 s (*c*) 3 s

Fig. 19. Performance of planning module (2)

2) **Ablation Analysis**

In the ablation analysis, we compare HybridLoss with three ablation variations: AblationLoss 1 (with s_f only), AblationLoss 2 (with s_p only), and AblationLoss 3 (without s_f and s_p). Including either s_f or s_p enhances the performance of both the prediction and planning modules, either individually or in combination with both s_f and s_p (Tables IX & X)

TABLE IX
ABLATION ANALYSIS (PREDICTION MODULE)

Group	s_f	s_p	ADE↓ (m)	FDE↓ (m)	MR↓	EPA↑
HybridLoss*	√	√	**1.11±0.03**	**1.36±0.03**	**0.14±0.00**	**0.28±0.00**
AblationLoss1	√		1.23±0.02	1.50±0.02	0.15±0.00	**0.28±0.00**
AblationLoss2		√	1.24±0.01	1.51±0.01	0.15±0.00	**0.28±0.00**
AblationLoss3			1.36±0.04	1.64±0.05	0.22±0.00	0.18±0.00

TABLE X
ABLATION ANALYSIS (PLANNING MODULE)

Group	s_f	s_p	L2↓ (m)			Col. Rate↓ (%)		
			1 *s*	2 *s*	3 *s*	1 *s*	2 *s*	3 *s*
HybridLoss*	√	√	**0.41±0.01**	**0.85±0.02**	**1.33±0.04**	**0.07±0.01**	0.13±0.03	**0.14±0.04**
AblationLoss1	√		0.53±0.08	1.11±0.05	1.55±0.10	0.08±0.03	**0.12±0.01**	0.17±0.02
AblationLoss2		√	0.49±0.08	0.97±0.13	1.37±0.11	0.08±0.03	0.16±0.04	0.20±0.02
AblationLoss3			0.55±0.08	1.04±0.14	1.59±0.25	0.11±0.03	0.21±0.03	0.23±0.04

3) **Robustness Analysis**

Table XI shows robustness under increasing perturbation levels (σ) for five objectives by two stability metrics, ΔPlan and Endpoint-Std. First of all, HybridLoss achieves the lowest ΔPlan (from 0.0235 m to 0.1384 m) and lowest Endpoint-Std (from 0.0188 m to 0.1118 m), while Loss1, Loss2, and Loss3 exhibit larger deviations that increase more rapidly as noise grows. Desensitized-Reg baseline is the second-best methods but it is still worse at every σ on both metrics. These results indicates that perturbation setting is effectively stressing planning stability, and widening gaps at higher σ highlight differences in sensitivity across losses. HybridLoss's consistent advantages over other methods suggest that substituting an adaptive, planning-oriented target during training improves robustness rather than relying on an explicit desensitization constraint. Overall, HybridLoss delivers the most stable motion planning under perturbations, strengthening the novelty claim as an adaptive robustness mechanism.

TABLE XI
ROBUSTNESS OF MOTION-PLANNING

(*a*) ΔPlan (m) / ↓

σ	Loss1	Loss2	Loss3	Desensitized-Reg	HybridLoss
0.05	0.0251	0.0294	0.0336	0.0237	**0.0235**
0.10	0.0499	0.0582	0.0661	0.0472	**0.0465**
0.20	0.0993	0.1162	0.1313	0.0945	**0.0921**
0.30	0.1490	0.1744	0.1965	0.1415	**0.1384**

(*b*) Endpoint Std (m) / ↓

σ	Loss1	Loss2	Loss3	Desensitized-Reg	HybridLoss
0.05	0.0202	0.0241	0.0389	0.0193	**0.0188**
0.10	0.0400	0.0479	0.0568	0.0384	**0.0374**
0.20	0.0799	0.0955	0.1133	0.0768	**0.0744**
0.30	0.1198	0.1431	0.1697	0.1151	**0.1118**

The boxplot at σ = 0.20 m shows that HybridLoss attains the lowest median ΔPlan with a relatively compact spread, while Loss3 exhibits the highest median deviation and largest overall displacement. Figs. 20-22 indicate that both ΔPlan and Endpoint-Std increase with perturbation level for all methods, but HybridLoss and Desensitized-Reg remain the lowest across full σ range. In contrast, Loss2 and Loss3 grow more

steeply as σ increases. These results suggest that HybridLoss not only reduces the average sensitivity of planned trajectories to input noise but also improves robustness across samples. Moreover, HybridLoss slightly but consistently outperforms the Desensitized-Reg baseline supports the argument that its adaptive loss design implicitly achieves stability without relying on hand-crafted robustness constraints. In summary, HybridLoss delivers the most stable planning under input perturbations, reinforcing its novelty relative to conventional trajectory optimization objectives.

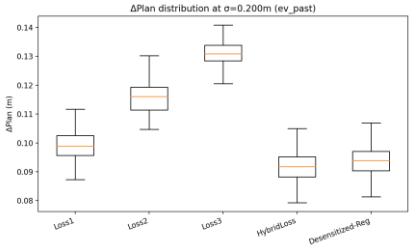

Fig. 20. ΔPlan distribution

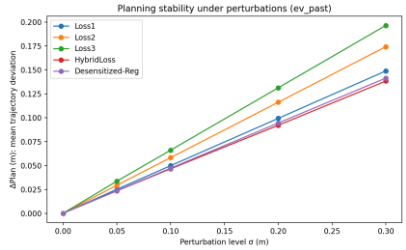

Fig. 21. Planning stability under perturbations for ΔPlan

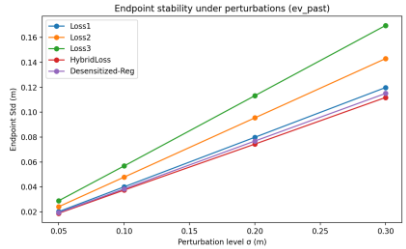

Fig. 22. Planning stability under perturbations for endpoint std

4) **RL Comparison (Generalization, Efficiency, Stability)**

Table XII shows that HybridLoss keeps almost the same success rate from seen to unseen conditions (87.9% vs 87.7%) while collision rate decreases from 1.42% to 1.23%. Trajectory accuracy slightly worsens on unseen data with ADE increasing from 0.599 m to 0.615 m and FDE from 1.078 m to 1.115 m. For computational demands, HybridLoss achieves 15.56 ms latency, 64.3 Hz, and only 0.87 M parameters, and its latency is lower than RL-QPSO and TC-DDPG although it is higher than Standard DDPG (Fig. 23). These results show that overall task success is stable and safety improves in unseen scenarios, suggesting that learned planning is not overfitting the training distribution. Meanwhile, the modest increases in ADE, FDE, lateral-acceleration RMS with a small rise in constraint violations, imply reduced smoothness and comfort under novel conditions but with limited degradation. In summary, HybridLoss offers a trade-off among adaptability, efficiency, and stability, remaining robust under unseen conditions with only minor performance drift.

TABLE XII
RL COMPARISON FOR ADAPTABILITY, EFFICIENCY, AND STABILITY

(*a*) Adaptability to unseen conditions

Indicator	Train (Seen)	Test (Unseen)	GAP (%)
Success Rate ↑	87.9%	87.7%	-0.5%
Collision Rate ↓	1.42%	1.23%	-12.7%
ADE (m) ↓	0.599	0.615	+2.6%
FDE (m) ↓	1.078	1.115	+3.4%

(*b*) Computational demands

Model	Latency (Mean)	Latency (P99)	FPS	Model params
HybridLoss	15.56 ms	16.75 ms	64.3 Hz	0.87 M

(*c*) Stability analysis

Indicator	Train (Seen)	Test (Unseen)	GAP (%)
Lat. Accel RMS	0.68	0.72	+6.1%
Action Smoothness	0.337	0.357	--
Constraint Viol.	5.2%	5.5%	--

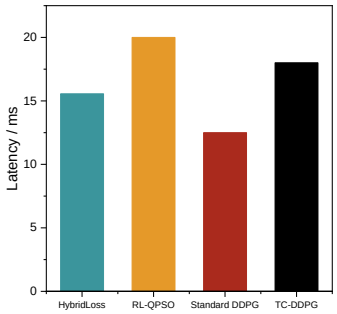

Fig. 23. Latency comparison

5) **Multi-Objective Trade-off Analysis of HybridLoss**

Across the five methods, the table shows efficiency (Avg. Speed), safety (MinDist), comfort (1/jerk), and a normalized overall score, while the scatter plot visualizes the efficiency-safety trade-off with point size indicating comfort. HybridLoss achieves the best overall score (0.5810) and the strongest combined safety and comfort while maintaining competitive efficiency with 5.2148 m/s. However, Loss3 only attains the highest effiency but collapses in safety and comfort, and Desensitized-Reg is relatively balanced but remain slower and less comfortable than HybridLoss with a decent overall score (0.5059). The plot places HybridLoss in the upper-right region with a larger marker, indicating a more favorable efficiency-safety balance without sacrificing ride quality. The results suggest that HybridLoss is not the merely reweighting objectives, it improves the Pareto-style trade-off by avoiding the speed-driven failure mode seen in Loss3 and by improving both efficiency and comfort over Desensitized-Reg while preserving strong safety. In summary, HybridLoss delivers the most balanced multi-objective planning performance among all compared approaches (Table XIII & Fig. 24).

TABLE XIII
MULTI-OBJECTIVE SUMMARY

Methods	Efficiency (Speed, m/s ↑)	Safety (MinDist, m ↑)	Comfort (1/Jerk, s^3/m, ↑)	Overall Score (After norm.)
HybridLoss	5.2148	7.3126	8.6949	0.5810
Desensitized-Reg	3.6772	6.5872	4.0247	0.5059
Loss1	3.8836	5.0909	6.2682	0.3202
Loss2	5.1681	5.9462	7.1931	0.4593
Loss3	9.3739	2.1737	2.1555	0.5000

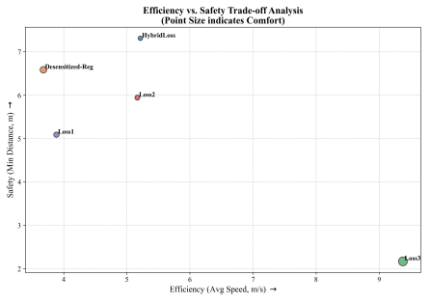

Fig. 24. Efficiency vs safety trade-off analysis

6) **Social Force Cost Analysis**

Fig. 25 and Table XIV show that social force cost distributions different methods. HybridLoss concentrates most probability at relatively low social cost, while Loss 2 and Desensitized-Reg follow a similar low-cost pattern with different peaks and spread. Loss 1 is shifted to the right compared with these methods, indicating a higher social cost, and Loss3 is the most right-shifted with a broad tail toward large costs, indicating the weakest implicit cooperation. Since lower social cost corresponds to more cooperative and less conflict-prone interactions, left-shift of HybridLoss suggests it encourages more interaction-aware, implicitly cooperative behavior than standard loss baseline. In summary, HybridLoss yields a more favorable cooperation profile by reducing social-force penalties, supporting its role as an effective approximation of cooperative behavior in mixed-traffic.

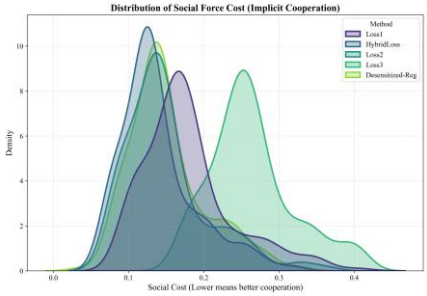

Fig. 25. Distribution of social force cost

TABLE XIV
MEAN VALUE OF SOCIAL FORCE COST

	Loss1	Loss2	Loss3	Desensitized-Reg	HybridLoss
Social force cost	0.1785	0.1495	0.2631	0.1512	0.1386

For further analysis of social force cost, Table XV compares five methods on conflict index, yield rate, minimum inter-agent distance, and collision rate, showing that HybridLoss achieves a relatively low conflict index (0.0196) and highest yield rate (0.0318) while maintaining a high minimum distance (29.0994 m) and one of the lowest collision rates (0.0030). In contrast, Loss3 has the lowest conflict index but also the smallest minimum distance, whereas Loss1 exhibits the highest collision rate despite a large minimum distance. Overall, HybridLoss provides the most balanced interaction behavior by combining stronger yielding with low collision risk and large safety buffers, indicating improved implicit cooperation in mixed-traffic scenarios. In summary, these results support that HybridLoss better approximates cooperative planning by reducing risky interactions while encouraging more socially compliant maneuvers.

TABLE XV
COMPARISON OF CONFLICT INDEX, YIELD RATE, MIN DIST, AND COLLISION RATE

Method	Conflict index ↓	Yield rate ↑	Min distance ↑	Collision rate ↓
Loss1	0.0223	0.0243	29.5921	0.0092
Loss2	0.0216	0.0104	28.8778	0.0087
Loss3	0.0164	0.0238	23.0874	0.0031
Desensitized-Reg	0.0212	0.0114	28.7270	0.0066
HybridLoss	0.0196	0.0318	29.0994	0.0030

7) **Constrained Maneuver Evaluation**

Across constrained maneuvers, HybridLoss gets the best overall performance with the lowest endpoint error, highest smoothness score, and largest minimum distance. In contrast, Loss3 shows the worst on motion quality with highest mean jerk and lowest smoothness score. These results show that HybridLoss not only improves high accuracy but also better maintains safety margins under tight constraints. Additionally, although its mean jerk is not the lowest one, HybridLoss provides a stronger balanced trade-off, avoiding the instability seen in Loss3 and the reduced safety margin of Loss2. In summary, the constrained-maneuver benchmark indicates that HybridLoss delivers the most reliable accuracy-smoothness-safety balance among all compared methods (Table XVI & Fig. 26).

TABLE XVI
CONSTRAINED MANEUVER EVALUATION

Methods	Endpoint error (m, ↓)	Mean jerk (m/s^3, ↓)	Smoothness score (↑)	Min distance (m, ↓)
Desensitized-Reg	4.5613	**15.0375**	0.0233	33.6870
Loss1	4.8708	22.3838	0.0173	35.8883
Loss2	5.2088	14.1464	0.0240	25.4030
Loss3	6.2171	57.7133	0.0076	25.8170
HybridLoss	**3.9720**	19.4235	**0.0262**	**37.3297**

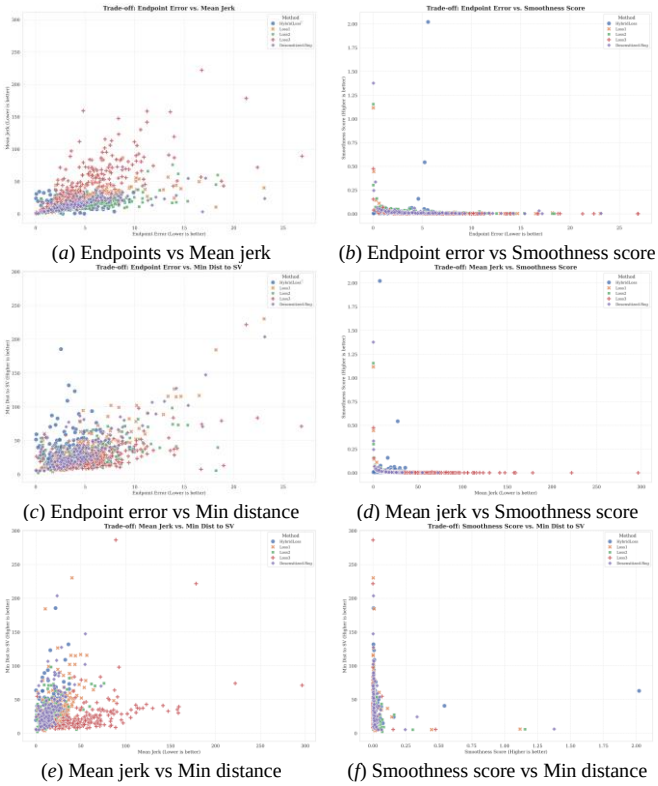

(*a*) Endpoints vs Mean jerk (*b*) Endpoint error vs Smoothness score

(*c*) Endpoint error vs Min distance (*d*) Mean jerk vs Smoothness score

(*e*) Mean jerk vs Min distance (*f*) Smoothness score vs Min distance

Fig. 26. Trade-off for constrained maneuver evaluation

V. Conclusion

This paper introduces HybridLoss, an adaptive and planning-oriented loss function designed to strengthen E2E prediction and planning for AVs. HybridLoss integrates two primary components: an adaptive motion-planning mechanism and a comprehensive composite training objective. Specifically, the adaptive motion-planning mechanism replaces the raw demonstrated AV trajectory in the planning loss with a scenario-adaptive, planner-generated reference, enabling the model to learn motion-planning strategies that are more consistent with rational planning objectives and less sensitive to suboptimal or overly aggressive demonstrations. This adaptive approach consists of three key elements: parameter definition, a planning optimization module, and a planning selection module.

HybridLoss objective combines multiple complementary terms, including prediction loss, planning loss (with adaptive motion planning), safety potential value loss, and social force value loss. By integrating these elements, HybridLoss tightens the coupling between prediction and planning and encourages a more effective closed-loop feedback mechanism between the two modules. As a result, HybridLoss promotes a balanced consideration of individual feasibility and interaction-aware behavior, making it particularly suitable for complex traffic scenarios. Its adaptive planning selection further improves robustness in mixed-traffic environments by dynamically choosing planning strategies that better match scenario characteristics [34]. Grounded in rational motion-planning principles, HybridLoss consistently outperforms conventional loss designs in our evaluations, demonstrating its effectiveness for improving E2E planning quality in autonomous driving systems.

While our current experiments primarily focus on controlled comparisons to isolate the contribution of the proposed loss design, we acknowledge the importance of benchmarking against broader state-of-the-art (SOTA) systems. In future work, we will integrate HybridLoss into representative recent SOTA prediction–planning frameworks and evaluate it on established public benchmarks, enabling a more comprehensive comparison under diverse settings and further validating the generality and practical impact of the proposed objective.

References

[1] P Bhavsar, P Das, M Paugh, K Dey, M Chowdhury, "Risk analysis of autonomous vehicles in mixed traffic streams," *Transportation Research Record*, vol. 2625, no. 1, pp. 51-61, Jan. 2017. doi: 10.3141/2625-06.

[2] L Le Mero, D Yi, M Dianati, A Mouzakitis, "A survey on imitation learning techniques for end-to-end autonomous vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 9, pp. 14128-14147, Sept. 2022. doi: 10.1109/TITS.2022.3144867.

[3] A Gupta, J Johnson, L Fei-Fei, S Savarese, A Alahi, "Social GAN: socially acceptable trajectories with generative adversarial networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Salt Lake City, Utah, USA, 2018, pp. 2255-2264.

[4] T Salzmann, B Ivanovic, P Chakravarty, M Pavone, "Trajectron++: dynamically-feasible trajectory forecasting with heterogeneous data," In *European Conference on Computer Vision*, Glasgow, UK, 2020, pp. 683-700.

[5] M Liang et al., "Learning lane graph representations for motion forecasting," In *European Conference on Computer Vision*, Glasgow, pp. 541-556.

[6] J Gao et al., "VectorNet: encoding hd maps and agent dynamics from vectorized representation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, USA, 2020, pp. 11525-11533.

[7] Y Chai, B Sapp, M Bansal, D Anguelov, "MultiPath: multiple probabilistic anchor trajectory hypotheses for behavior prediction," *arXiv preprint arXiv:1910.05449*, pp. 1-14.

[8] T Phan-Minh, EC Grigore, FA Boulton, O Beijbom, EM Wolff, "CoverNet: multimodal behavior prediction using trajectory sets," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, USA, 2020, pp. 14074-14083.

[9] H Zhao et al., "TNT: target-driven trajectory prediction," in *Proceedings of the 2020 Conference on Robot Learning, PMLR*, Boston, USA, 2020, pp. 895-904.

[10] J Gu, C Sun, H Zhao, "DenseTNT: end-to-end trajectory prediction from dense goal sets," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, Montreal, QC, Canada, 2021, pp. 15303-15312.

[11] S Shi, L Jiang, D Dai, B Schiele, "Motion transformer with global intention localization and local movement refinement," in *Advances in Neural Information Processing Systems*, New Orleans, Louisiana, USA, 2022, pp. 6531-6543.

[12] N Nayakanti, R Al-Rfou, A Zhou, A Goel, J Sacks, "Wayformer: motion forecasting via simple & efficient attention networks," *arXiv preprint arXiv:2207.05844*, pp. 1-20.

[13] C Bai, X Gao, M. Chen, W Guo, D Rong, C Yang, and S Jin, "Hybrid data-model driven trajectory prediction on highways: integrating anticipatory interaction awareness and personalized driving preferences," *Transportation Research Part C: Emerging Technologies*, vol. 180, pp. 105351, Nov. 2025, doi: 10.1016/j.trc.2025.105351.

[14] S Ross, G Gordon, JA Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, Fort Lauderdale, FL, USA, 2011, pp. 627-635.

[15] M Bansal, A Krizhevsky, A Ogale, "ChauffeurNet: learning to drive by imitating the best and synthesizing the worst," *arXiv preprint arXiv:1812.03079*, pp. 1-20.

[16] C Bai, C Yang, D Rong, W Guo, X Gao, W Yao, S Jin, "Multi-source temporal attention fusion network (MTAFN) for driving risk assessment based on naturalistic driving data," *Expert Systems with Applications*, vol. 279, pp. 127502, Jun. 2025, doi: 10.1016/j.eswa.2025.127502.

[17] D Chen, B Zhou, V Koltun, P Krähenbühl, "Learning by Cheating," in *Proceedings of the Conference on Robot Learning, Cambridge*, MA, USA, 2020, pp. 66-75.

[18] K Renz et al., "PlanT: explainable planning transformers via object-level representations," *arXiv preprint arXiv:2210.14222*, pp. 1-12.

[19] D Rong et al., "Smart prediction-planning algorithm for connected and autonomous vehicle based on social value orientation," J*ournal of Intelligent and Connected Vehicles*, vol. 8, no. 1, pp. 1-17, Mar. 2025. doi: 10.26599/JICV.2024.9210053.

[20] D Rong et al., "Integration of multi-vehicle prediction and planning based on social value orientation in mixed traffic," *IEEE Transactions on Intelligent Vehicles*, pp. 1-13, Aug. 2024. doi: 10.1109/TIV.2024.3442214.

[21] D Rong et al., "Multi-vehicle collaborative trajectory planning based on Kaldor-Hicks improvements," *IEEE Transactions on Automation Science and Engineering*, vol. 22, pp. 12512-12525, Feb. 2025, doi: 10.1109/TASE.2025.3542840.

[22] B Amos, JZ Kolter, "OptNet: differentiable optimization as a layer in neural networks," in *International Conference on Machine Learning*, Sydney, Australia, 2017, pp. 136-145.

[23] B Amos et al., "Differentiable MPC for End-to-end Planning and Control," in *Advances in neural information processing systems*, Montréal, Canada, 2018, pp. 31.

[24] Z Huang, H Liu, J Wu, C Lv, "Differentiable integrated motion prediction and planning with learnable cost function for autonomous driving," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 35, no. 11, pp. 15222-15236, Nov. 2024. doi: 10.1109/TNNLS.2023.3283542.

[25] H Liu et al., "Hybrid-prediction integrated planning for autonomous driving," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 47, no. 4, pp. 2597-2614, April 2025. doi: 10.1109/TPAMI.2025.3526936.

[26] W Zeng et al., "End-to-end interpretable neural motion planner," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Long Beach, CA, USA, 2019, pp. 8660-8669.
[27] Y Hu et al., "Planning-oriented autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Vancouver, BC, Canada, 2023, pp. 17853-17862.
[28] S Hu et al., "St-p3: End-to-end vision-based autonomous driving via spatial-temporal feature learning," in *European Conference on Computer Vision*, Milan, Italy, 2022, pp. 533-549.
[29] Y Yan et al., "A multi-vehicle game-theoretic framework for decision making and planning of autonomous vehicles in mixed traffic," *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 11, pp. 4572-4587, Nov. 2023. doi: 10.1109/TIV.2023.3321346.
[30] H J Guan, B Y Wang, J W Gong, and H Y Chen, "Coordinated motion planning for heterogeneous autonomous vehicles based on driving behavior primitives," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 11, pp. 11934-11949, Nov. 2023. doi: 10.1109/TITS.2023.3285904.
[31] S Mohammadian, Z Zheng, M Haque, and A Bhaskar, "NET-RAT: Non-equilibrium traffic model based on risk allostasis theory," *Transportation Research Part A: Policy and Practice*, vol. 174, pp. 1-25, Aug. 2023. doi: 10.1016/j.tra.2023.103731.
[32] J L Kong, X M Fan, X B Jin, S Lin, and M. Zuo, "A variational bayesian inference-based En-Decoder framework for traffic flow prediction," *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 3, pp. 2966-2975, Mar. 2024. doi: 10.1109/TITS.2023.3276216.
[33] R Chai, T Antonios, S Chai, Y Xia, S Al, CL Chen, "Multiphase overtaking maneuver planning for autonomous ground vehicles via a desensitized trajectory optimization approach," *IEEE Transactions on Industrial Informatics*, vol. 19, no. 1, pp. 74-87, Apr. 2022. doi: 10.1109/TII.2022.3168434.
[34] W Yao et al., "Unravelling heterogeneity of commuters' travel behavior: an empirical investigation of commuting regularity using license plate recognition data," *Transportation*, pp. 1-37, Dec. 2025, doi: 10.1007/s11116-025-10713-7.

Donglei Rong is currently a Ph.D. candidate in the Institute of Intelligent Transportation Systems, College of Civil Engineering and Architecture, Zhejiang University and The Hong Kong Polytechnic University. His research interests include intelligent transportation systems, autonomous vehicle, trajectory planning, and traffic modeling.

Chengcheng Yang received the B.S. degree from the School of Automotive and Traffic Engineering at Jiangsu University, Zhenjiang, Jiangsu, China, in 2017; the M.S. degree from the School of management at the University of Shanghai for Science and Technology, Shanghai, China, in 2021. He is currently working toward the Ph.D degree with the Institute of Intelligent Transportation Systems of the College of Civil Engineering and Architecture, Zhejiang University, Hangzhou, Zhejiang, China. His research interest including traffic flow theory and traffic control.

Congcong Bai recieved the B.S. degree in the Institute of Intelligent Transportation Systems of the College of Civil Engineering and Architecture, Zhejiang University, Hangzhou, Zhejiang, China, 2021. He is currently working toward the Ph.D degree with the Institute of Intelligent Transportation Systems of the College of Civil Engineering and Architecture, Zhejiang University, Hangzhou, Zhejiang, China. His research interest including micro driving behavior modeling, automatic driving behavior prediction and traffic big data mining.

Wentong Guo recieved the B.S. degree in the College of Water Conservancy and Engineering, Zhengzhou University, Zhengzhou 450001, China, 2022. He is currently working toward the M.S. degree with the Institute of Intelligent Transportation Systems of the College of Civil Engineering and Architecture, Zhejiang University, Hangzhou, Zhejiang, China. His research interest including micro driving behavior modeling, built environment analysis and traffic big data mining.

Sheng Jin received the Ph.D. degree in Transportation Engineering from Jilin University, Changchun, China, in 2010. Now, he is a Professor with the College of Civil Engineering and Architecture, Zhejiang University, Hangzhou, China. He has published more than 80 articles in related journals such as Transportation Research part A/B/C/D, Accident Analysis & Prevention, IEEE T ITS, and IEEE ITSM. His research is focused on traffic flow theory, intelligent transportation systems, traffic signal control, and traffic big data.

Min Xu received the bachelor's degree in hydraulic engineering and economics (double degree) from Tsinghua University and the Ph.D. degree in transportation engineering from the National University of Singapore, Singapore. She is currently an Associate Professor with the Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University. Her research interests include modeling and optimization problem in urban transportation and logistics systems with emerging technologies. She serves on the editorial advisory board for *Transportation Research Part C and Transportation Research Part E*.