

ExpertAD: Enhancing Autonomous Driving Systems with Mixture of Experts

Haowen Jiang, Xinyu Huang, You Lu, Dingji Wang, Yuheng Cao, Chaofeng Sha, Bihuan Chen*, Keyu Chen, Xin Peng

College of Computer Science and Artificial Intelligence, Fudan University, Shanghai, China

Abstract

Recent advancements in end-to-end autonomous driving systems (ADSs) underscore their potential for perception and planning capabilities. However, challenges remain. Complex driving scenarios contain rich semantic information, yet ambiguous or noisy semantics can compromise decision reliability, while interference between multiple driving tasks may hinder optimal planning. Furthermore, prolonged inference latency slows decision-making, increasing the risk of unsafe driving behaviors. To address these challenges, we propose ExpertAD, a novel framework that enhances the performance of ADS with Mixture of Experts (MoE) architecture. We introduce a Perception Adapter (PA) to amplify task-critical features, ensuring contextually relevant scene understanding, and a Mixture of Sparse Experts (MoSE) to minimize task interference during prediction, allowing for effective and efficient planning. Our experiments show that ExpertAD reduces average collision rates by up to 20% and inference latency by 25% compared to prior methods. We further evaluate its multi-skill planning capabilities in rare scenarios (e.g., accidents, yielding to emergency vehicles) and demonstrate strong generalization to unseen urban environments. Additionally, we present a case study that illustrates its decision-making process in complex driving scenarios.

Extended version — <https://arxiv.org/abs/2511.11740>

Introduction

End-to-end autonomous driving systems (ADSs) are gaining growing attention, supported by advances in computational power that enable efficient data processing and real-time decisions. Most vision-based ADSs (Hu et al. 2022; Jiang et al. 2023) use multi-view camera inputs to produce 2D representations for downstream perception, prediction, and planning. Recent efforts integrate large language models (LLMs)(Shao et al. 2024; Xu et al. 2024) and world models (WMs)(Wang et al. 2025b, 2023), leveraging their generalization and predictive capabilities. However, most deployed systems (Tian et al. 2024; Chen et al. 2025) still rely on traditional end-to-end ADSs to generate final trajectories. Even with LLMs or WMs, end-to-end ADSs are typically retained as fallbacks to ensure reliable execution. In this work, we

*Corresponding Author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

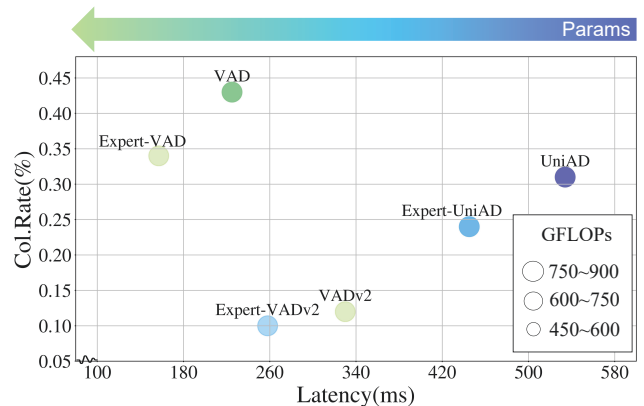


Figure 1: **Collision rate and latency trade-offs across different models.** ExpertAD exhibits substantial improvements in planning effectiveness while reducing inference latency, as measured on NVIDIA GeForce RTX 3090.

aim to enhance end-to-end ADSs by improving planning effectiveness and reducing inference latency.

One key challenge is to deal with ambiguous semantic information during inference, which can hinder reliable decision-making. For instance, critical environmental features may be partially captured due to sensor noise or occlusion, leading to gaps in contextual understanding. Therefore, recalibrating feature channels could help ADS reconstruct missing context. Another important challenge lies in the varying relevance of driving tasks across different scenarios. For instance, mapping task aids curved-road planning but is less useful for straight paths. This suggests that fully activating all tasks is unnecessary, as different scenarios demand different tasks. Therefore, dynamic task selection might minimize the interference between driving tasks, achieving optimal planning and reduce inference latency.

To selectively amplify task-critical features and dynamically choose relevant tasks, we introduce ExpertAD, a novel framework that integrates the lightweight efficiency of the Mixture of Experts (MoE) paradigm (Shazeer et al. 2017) into ADSs. While Mixture-of-Experts (MoEs) have shown promise across domains, their adoption in autonomous driving remains limited. Particularly, minor input variations in

dynamic scenes can destabilize expert activation, and most prior works apply MoEs solely within the planning module, often using shared low-rank experts without task-specific specialization. For instance, ARTEMIS (Feng et al. 2025) incorporates MoE into an autoregressive planner to improve temporal modeling, but this comes at the cost of significant inference latency due to sequential token-wise processing.

In contrast, ExpertAD achieves both effectiveness and efficiency through an end-to-end MoE architecture that spans perception and planning. A Perception Adapter (PA) dynamically reweights BEV features to emphasize task-relevant semantics (e.g., pedestrians, road geometry), while a Mixture of Sparse Experts (MoSE) enables efficient and stable expert activation based on the driving context. MoSE customizes expert realization using sparse attention over long-term historical features, reducing computation while preserving behavioral diversity. Our framework supports task adaptivity, enhances generalization to unseen scenarios, and enables low-latency decision making, which offers a superior trade-off compared to planning-only MoE frameworks.

We integrate the ExpertAD framework with three state-of-the-art vision-only ADSs (*i.e.*, UniAD (Hu et al. 2023), VAD (Jiang et al. 2023) and VADv2 (Chen et al. 2024)) and conduct experiments to evaluate the effectiveness of ExpertAD. As depicted in Fig. 1, ADSs with our ExpertAD framework achieve reductions in collision rates by about 20% while reducing inference latency by approximately 25%. Moreover, ExpertAD shows improved multi-skill planning capabilities in certain scenarios. We also conduct a new-city generalization study to validate the generalization capability of ExpertAD and a case study to visualize the expert pathways engaged in decision making.

This work makes the following main contributions.

- We propose a Perception Adapter (PA) to amplify task-critical features and a Mixture of Sparse Experts (MoSE) to reduce the driving task interference.
- We propose ExpertAD, which is a novel framework that integrates the lightweight efficiency of the Mixture of Experts (MoE) framework into ADSs.
- We integrate ExpertAD into three state-of-the-art vision-only ADSs, and conduct large scale experiments to show improved planning effectiveness and lower inference latency across both open-loop and closed-loop datasets.

Preliminary and Related Work

We first introduce the preliminary on modular end-to-end ADSs, and then discuss and review the relevant work.

Preliminary

In a typical modular end-to-end ADS pipeline, sequential modules work together for perception, prediction and planning. First, multi-view images are processed by a BEV encoder (Prakash, Chitta, and Geiger 2021; Li et al. 2022a; Liang et al. 2022) to generate BEV features, which are essential for understanding the environment. The perception module then utilizes these BEV features for tasks like multi-object tracking (Zeng et al. 2022; Zhang et al. 2022; Li and Jin 2022) and online mapping (Kirillov et al. 2019;

Kim et al. 2020; Li et al. 2022b), providing a precise understanding of the surroundings. Next, the prediction module performs tasks like ego-state estimation, environmental interaction modeling, and navigation execution to forecast the future trajectories of ego vehicle and surrounding objects (*i.e.*, vehicles and pedestrians). Finally, the planning module (Casas, Sadat, and Urtasun 2021; Ye et al. 2023; Wang et al. 2021) synthesizes perceptual and predictive inputs to generate a collision-free trajectory.

Related Work

End-to-End Autonomous Driving. Early methods (Pomerleau 1988) use neural networks to process video and laser range data, generating steering commands via imitation learning (Xiao et al. 2023; Zhang et al. 2021) and reinforcement learning (Toromanoff, Wirbel, and Moutarde 2020). Recent works (Shao et al. 2023; Sun et al. 2023; Li et al. 2025; Jia et al. 2025) enhance modular end-to-end frameworks by incorporating intermediate tasks and advanced architectures for joint optimization. ST-P3 (Hu et al. 2022) integrates perception, prediction and planning modules for richer features, while UniAD (Hu et al. 2023) unifies these modules through a planning-oriented pipeline to eliminate accumulative errors. VAD (Jiang et al. 2023) employs vectorized representations for efficient planning, and VADv2 (Chen et al. 2024) introduces probabilistic modeling for safer planning. Recently, large language models and world models have attracted attention for their improved planning effectiveness and strong generalization capabilities. LLMs/MLLMs in autonomous driving (Shao et al. 2024; Xu et al. 2024; Sima et al. 2023) provide textual descriptions as well as control signals for diverse driving scenarios. World models (Wang et al. 2025b, 2024, 2023; Yang et al. 2023; Zhang et al. 2024b; Gao et al. 2023) build comprehensive environment representations and make planning decisions based on predicted future states. Despite their promise, high inference latency remains a key barrier to real-world deployment. Even when LLMs or world models are deployed, vehicles often retain a conventional end-to-end model (Tian et al. 2024; Wang et al. 2025a) as a fallback to ensure robust execution. To reduce inference time of end-to-end ADSs, DriveAdapter (Jia et al. 2023) adopts a teacher-student paradigm where the student model gains driving knowledge from the teacher model and aligns features between privileged and raw sensor inputs. PlanKD (Feng et al. 2024) further enables student planner to inherit planning-relevant knowledge to improve performance. However, these two approaches pursue the lower inference latency at the cost of planning effectiveness. Our work aims to improve both planning effectiveness and inference efficiency of modular end-to-end ADSs with a MoE architecture.

Mixture-of-Experts (MoE). MoE models aim to scale model capacity while reducing computational costs by activating only a subset of parameters for each input. Recent works utilize MoE in large language models (Du et al. 2022; Team 2023; Jiang et al. 2024) and vision models (Daxberger et al. 2023; Riquelme et al. 2021) to improve training efficiency and inference performance. In autonomous driving, some approaches (Ohn-Bar et al. 2020; Nazeri and Bohlouli

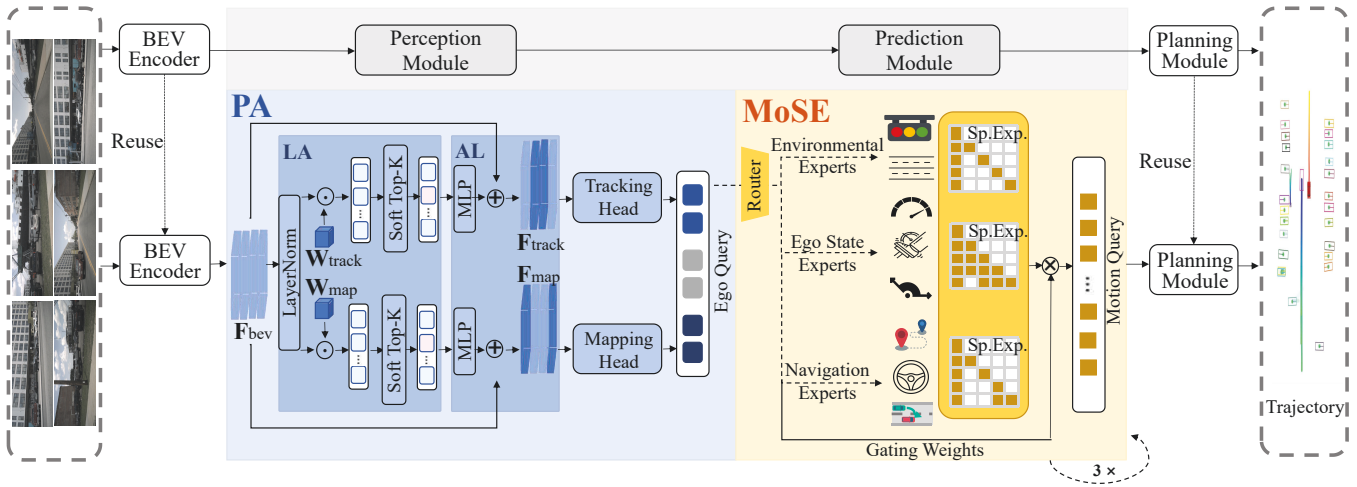


Figure 2: **Overall architecture of ExpertAD.** ExpertAD is built upon ADS models by retaining all the original modules except for the perception and prediction modules. The perception module is restructured as the Perception Adapter (PA) to amplify task-critical features, enhancing scene understanding. The prediction module is transformed into the Mixture of Sparse Experts (MoSE), minimizing interference among driving tasks and improving overall open-loop planning performance.

2021) rely on supervised learning from historical driving data, with MoE facilitating the execution of behavior instructions in predefined driving scenarios. Despite their potential, these models still have unsatisfactory planning effectiveness, facing challenges in generalization abilities to adjust their behavior across unseen driving scenarios. Some MoE-based approaches have been applied to the perception module, such as camera selection (Fang and Choromanska 2020; Morra et al. 2023) and visual token selection (Zhang et al. 2024a). Others focus on the planning module, targeting tasks like trajectory selection (Feng et al. 2025; Pini et al. 2023) and driving intention prediction (Yuan et al. 2023). While effective for specific subtasks, these methods often lack unified semantic and behavioral modeling, making them highly sensitive to noise in both perception and planning. Moreover, they typically introduce additional computational overhead during inference. In contrast, our proposed framework, ExpertAD, amplifies task-critical features and reduces task interference, optimizing both planning effectiveness and inference efficiency.

Methodology

We design and implement ExpertAD, a novel Mixture of Experts (MoE) framework, to improve the planning effectiveness and inference efficiency of ADSs.

Approach Overview

The framework overview of ExpertAD is shown in Fig. 2. ExpertAD is built upon an ADS, retaining all original modules except for the perception and prediction modules. We restructure the perception module as the **Perception Adapter (PA)** to amplify task-related features for context-aware scene understanding. We transform the prediction module into the **Mixture of Sparse Experts (MoSE)** to dynamically activate driving tasks, minimizing interference.

We define a combined **Training Loss** to ensure task effectiveness while promoting efficient expert utilization.

As multi-view camera images are encoded into BEV features through the BEV Encoder, the PA selects task-relevant feature dimensions via a *Learned Adapter (LA)*, and amplifies task-related semantics for the tracking and mapping transformers using *Alignment Layers (AL)*. The PA outputs an ego query, combining the agent query (*i.e.*, the output of tracking head), the map query (*i.e.*, the output of mapping head), and a learnable embedding. Prediction tasks are categorized into three groups with eight *Sparse Experts (Sp.Exp.)*, each using a sparse attention mechanism based on their respective focus. The MoSE uses the ego query from the PA to dynamically activate the most relevant experts through a *Router* based on the current driving scenario. The motion query output from the MoSE is then passed to the planning module to determine the final driving route.

Perception Adapter

BEV features from the BEV Encoder capture various semantics (*e.g.*, roads, vehicles, and traffic signs). However, the focus of each perception task (*i.e.*, multi-object tracking (Zeng et al. 2022; Zhang et al. 2022) and online mapping (Kirillov et al. 2019; Kim et al. 2020; Li et al. 2022b)) differs. Directly passing all features to the tracking and mapping transformers might cause non-critical dimensions to overshadow key features, leading to information loss. To address this, we propose PA, composed of a *Learned Adapter* and an *Alignment Layer*, to amplify task-critical features.

Learned Adapter. It is designed to select a subset of task-specific BEV features for each task in the perception module. Instead of relying on static feature selection methods that remove or merge features based on predefined criteria (Bolya et al. 2022; Rao et al. 2021; Yin et al. 2022), we adapt the idea of conditional adapter (Lei et al. 2023), originally designed for token selection in fine-tuning.

Specifically, We first normalize and pool the BEV features along the temporal dimensions to obtain a frame-agnostic BEV representation $\tilde{\text{BEV}} \in \mathbb{R}^{d \times H \times W}$, where d is the feature dimension and $H \times W$ is the spatial resolution from the BEV Encoder. For each task $t \in T$ (e.g., tracking, mapping), we introduce a task-specific learnable parameter $w^{(t)} \in \mathbb{R}^d$ to dynamically weight feature dimensions. The initial task-aware selection score $s^{(t)} \in \mathbb{R}^d$ is computed by Eq. 1, where $\tilde{\text{BEV}}_{:,i,j}$ denotes the d -dimensional feature vector at spatial position (i, j) . This score reflects the importance of each feature dimension for the given task.

$$s = \frac{1}{H \times W} \sum_{i,j} \tilde{\text{BEV}}_{:,i,j} \odot w \quad (1)$$

Based on the task-specific score $s^{(t)}$, we compute the soft channel selection weights $\lambda^{(t)} \in [0, 1]^d$ by solving the constrained optimization problem defined in Eq. 2,

$$\begin{aligned} f(s) &:= \arg \max_{\lambda} (s)^\top \lambda + \epsilon \Omega(\lambda) \\ \text{s.t. } \mathbf{1}^\top \lambda &= \tau, \lambda \in [0, 1]^d \end{aligned} \quad (2)$$

where the entropy $\Omega(\lambda)$ regularizes the solution of λ , promoting smoothness in the feature selection process. The constant τ represents the number of key feature dimensions. The constraint $\mathbf{1}^\top \lambda = \tau$ ensures the model focuses on τ dominant channels and the open-interval constraint $\lambda \in [0, 1]^d$ enables continuous weighting of channels. Since this optimization problem does not have a closed-form solution for $\epsilon > 0$ and $\tau > 1$, we employ the iterative algorithm by Lei *et al.* (Lei *et al.* 2023) to get subdifferentiable soft weights for λ .

Alignment Layer. It is designed to reweight feature dimensions to recalibrate and amplify task-critical features.

Specifically, given the selection weight $\lambda^{(t)}$ from the *Learned Adapter*, we obtain the calibrated features F_{align} by Eq. 3,

$$F_{align} = MLP(\text{BEV} \odot \lambda) + \text{BEV} \quad (3)$$

This equation combines BEV features from the BEV Encoder with amplified features. The term $\text{BEV} \odot \lambda$ represents element-wise multiplication, effectively amplifying certain features based on the learned weight distribution. The *MLP* layer further introduces non-linearity to the transformed features. The result is added back to the original BEV features to generate F_{align} , which not only preserves the original spatial information but also creates a shortcut path for gradients during backpropagation.

The aligned features F_{align} are then processed by the tracking transformer and mapping transformer to generate agent and map queries, providing a comprehensive understanding of the surroundings. Finally, the two queries are concatenated with a learnable embedding to form an ego query $\mathcal{F}_{ego} \in \mathbb{R}^{B \times L \times d}$, where B is the batch size, L is the sequence length and d is the feature dimension. This ego query serves as the input to the MoSE, allowing the model to process the concatenated information for further tasks.

Mixture of Sparse Experts

The prediction module includes various driving tasks to enhance ADS decision-making. However, interactions among these tasks, especially with long historical sequences, can cause interference and raise computational cost. To address this, we propose MoSE, which uses *Sparse Experts* and a *Router* to dynamically activate tasks based on the ego query \mathcal{F}_{ego} from the previous module.

Sparse Experts. To forecast future trajectories of the ego vehicle and nearby agents, the prediction module coordinates three critical tasks: ego-state estimation, environmental interaction modeling and navigation execution. These are formalized into three expert categories, *i.e.*, *Environmental Experts*, *Ego State Experts* and *Navigation Experts*, with a total of eight experts.

Environmental Experts process environmental information provided by the perception module, such as traffic signals, lane markings, and obstacles. Specifically, the *Tracking Expert* focuses on modeling dynamic foreground objects, while the *Mapping Expert* maintains map topology with lane-graph connectivity.

Ego State Experts regulate vehicle dynamics based on historical trajectories. Specifically, the *Velocity Expert*, *Yaw Expert* and *Acceleration Expert* collaborate to process vehicle-specific information for smooth trajectories.

Navigation Experts execute hierarchical route planning for vehicles. Specifically, The *Reference Point Expert* aligns trajectories with reference points, the *BEV Expert* provides spatial context, and the *Command Expert* interprets and executes navigation commands to adjust vehicle behavior.

Each expert is equipped with a sparse attention mechanism $MHCA()$ defined by Eq. 4 based on respective focus,

$$MHCA(Q, K, V) = \sum_{j \in C_i} \text{softmax} \left(\frac{Q_i K_j^T}{\sqrt{d_k}} \right) V_j \quad (4)$$

where Q_i denotes the i -th query vector, K_j and V_j represent the selected keys and values, and C_i defines the sparse computation for the i -th query.

Specifically, *Environmental Experts* adopt block-wise sparse attention, *i.e.*, $C_i = \{j \mid \lfloor j/m \rfloor = \lfloor i/m \rfloor\}$, restricting attention to local regions with a block size of m , allowing efficient spatial pattern capture within environmental regions (e.g., obstacles and traffic signs) while reducing computational costs. *Ego State Experts* utilize a sliding window attention mechanism with a window size of w , where $C_i = [i - w, i + w]$, enabling rapid adaptation to changes in ego state (e.g., speed, and acceleration) while maintaining manageable computational complexity. *Navigation Experts* adopt a global TopK attention, defined by $C_i = \{j \mid K_j \in \text{Top}_k(Q_i K^T)\}$, selecting key positions with the highest attention scores for the current query to capture long-range dependencies and prioritize critical information across the entire scene, such as waypoints and navigation cues.

Each expert is tailored to the specific task required by the prediction module and encodes domain-specific data through a fully connected layer with modality-specific normalization. Specifically, navigation experts embed rout-

ing instructions as text features, environmental experts extract geometric features from perception outputs, and ego-state experts represent vehicle dynamics as time-series data. These embeddings, denoted as \mathcal{F}_{expert} , are fused with the ego query \mathcal{F}_{ego} from the Perception Adapter (PA) via the sparse attention mechanism $MHCA()$ (Eq. 5), yielding expert-informed queries $\bar{\mathcal{F}}_{expert}$. This modular design facilitates efficient integration of diverse inputs into a unified predictive framework.

$$\bar{\mathcal{F}}_{expert} = MHCA(\mathcal{F}_{ego}, \mathcal{F}_{expert}, \mathcal{F}_{expert}) \quad (5)$$

Router. Each expert in the prediction module generates a $\bar{\mathcal{F}}_{expert}$. We hope to activate different experts based on the various driving scenarios. The *Router* is responsible for weighting the top- k experts through logits to dynamically select experts rather than fully activating all experts.

To compute expert logits, we adopt the gating mechanism from MoE (Shazeer et al. 2017) and introduce a learnable parameter $\mathbf{W}_{gate} \in R^{d \times N}$, where d is the feature dimension and N denotes the number of experts, mapping \mathcal{F}_{ego} to relevant experts over a batch. To introduce stochasticity, Gaussian noise $\eta \sim \mathcal{N}(0, 1)$ is added to the logits during training, scaled by a learnable matrix \mathbf{W}_{noise} of the same shape as \mathbf{W}_{gate} . The logits are then computed using \mathcal{F}_{ego} by Eq. 6,

$$g(x) = x \cdot \mathbf{W}_{gate} + \eta \cdot (\text{softplus}(x \cdot \mathbf{W}_{noise}) + \epsilon) \quad (6)$$

$$\eta \sim \mathcal{N}(0, 1), \quad \epsilon > 0$$

where $\eta \sim \mathcal{N}(0, 1)$ and $\epsilon > 0$ is a positive factor controlling the noise scale. This noise-enhanced gating mechanism enables dynamic expert selection based on input features while incorporating randomness to prevent overfitting.

Then, we introduce a function $\mathcal{R}(\cdot)$, defined in Eq. 7, to select the top- k experts based on the logits of experts.

$$\mathcal{R}(x) = \text{Top}K(\text{softmax}(g(x)), k) \quad (7)$$

Taking \mathcal{F}_{ego} as input, the result of $\mathcal{R}(\mathcal{F}_{ego})$ represents the routing scores of the top- k experts. The gating function $\mathcal{R}(\cdot)$ employs a decoupled selection mechanism based on high-level scene understanding, avoiding information from experts that could introduce noise.

Finally, we compute the motion query \mathcal{F}_{Motion} of the selected top- k experts by Eq. 8,

$$\mathcal{F}_{Motion} = \sum_{i=1}^k \mathcal{R}(\mathcal{F}_{ego})_i \cdot \bar{\mathcal{F}}_{expert_i} \quad (8)$$

where $\mathcal{R}(\mathcal{F}_{ego})_i$ and $\bar{\mathcal{F}}_{expert_i}$ respectively represent the routing scores and the expert-specific query for the selected expert i . The resulting \mathcal{F}_{Motion} integrates and reweights ego features, environmental features and navigation features, providing a global feature representation tailored to the current driving scenario for planning.

Training Loss

Shazeer *et al.* (Shazeer et al. 2017) has observed that models tend to focus higher weights on a few experts, leading to underutilization of others and inefficient task distribution.

To address this, we incorporate a switch-loss (Fedus, Zoph, and Shazeer 2022) to encourage more balanced weight allocation across the experts.

Switch-loss balances the allocation of queries over a batch across N experts by penalizing experts that receive disproportionately high routing probabilities by Eq. 9,

$$\mathcal{L}_{switch} = N \cdot \sum_{i=1}^N f_i \cdot \mathcal{P}_i \quad (9)$$

where f_i is the actual query load on expert i , while \mathcal{P}_i is its expected routing probability. Switch-Loss penalizes experts that are overused relative to their expected share.

The final model loss, defined by Eq. 10, incorporates the losses from all tasks along with the switch loss,

$$\mathcal{L}_{total} = \alpha_1 \mathcal{L}_{perception} + \alpha_2 \mathcal{L}_{prediction} \quad (10)$$

$$+ \alpha_3 \mathcal{L}_{planning} + \alpha_4 \mathcal{L}_{switch},$$

where $\mathcal{L}_{perception}$, $\mathcal{L}_{prediction}$ and $\mathcal{L}_{planning}$ respectively represent the loss functions for the perception, prediction and planning modules in the baseline ADSs.

Experiments

We first introduce the implementation of our experiments. Then, we report the overall performance and modular performance of ExpertAD. Finally, we conduct an ablation study, a generalization evaluation and a qualitative study.

Implementation Details

Dataset. We evaluate ExpertAD using both open-loop and closed-loop datasets. For open-loop evaluation, we adopt the nuScenes dataset (Caesar et al. 2019), which contains 1,000 real-world scenes with sensor data from 6 cameras, 1 LiDAR, 5 radars, GPS, and IMU. We use only the 6 camera images as visual input. For closed-loop evaluation, we use the Bench2Drive benchmark (Jia et al. 2024) with 2 million annotated training frames from 12 towns and 23 weather conditions in CARLA (Dosovitskiy et al. 2017). Its evaluation set includes 220 routes across 44 interactive scenarios (denoted as Bench2Drive220).

Baselines. Our work is orthogonal to existing transformer based frameworks and can be integrated into any of them. To evaluate the effectiveness and efficiency of ExpertAD, we adopt three state-of-the-art models, *i.e.*, UniAD (Hu et al. 2023), VAD (Jiang et al. 2023), and VADv2 (Chen et al. 2024), as representative and publicly available baselines.

Training. We instantiate ExpertAD into UniAD, VAD and VADv2, denoted as Expert-UniAD, Expert-VAD, and Expert-VADv2. All approaches are trained using the same hyperparameters as respective baselines for fair comparison. Experiments are run on 8 NVIDIA Tesla A100 GPUs.

Metrics. Following previous works, we use ST-P3 (Hu et al. 2022) metrics, *i.e.*, L2 errors and collision rates, to evaluate open-loop driving effectiveness of ADSs. For closed-loop evaluation, we adopt Success Rate (SR), Route Completion (RC), and Driving Score (DS) (Dosovitskiy et al. 2017) to assess effectiveness. We further report Latency (average inference time per forward pass), GFLOPs

Approach	Open-loop Metric		Closed-loop Metric			Efficiency Metric		
	Avg.col ↓	Avg.L2 ↓	DS ↑	SR ↑	RC ↑	Latency ↓	GFLOPs ↓	Params ↓
UniAD	0.31	1.03	44.62	14.09	68.68	534 ± 18 ms	~ 856	~ 89M
Expert-UniAD	0.24	0.89	55.49	20.63	81.04	445 ± 20 ms	~ 728	~ 125M
VAD	0.43	1.21	43.31	17.27	61.60	225 ± 25 ms	~ 558	~ 58M
Expert-VAD	0.34	1.10	52.53	19.53	76.73	157 ± 23 ms	~ 461	~ 90M
VADv2	0.12	0.33	75.90	55.01	90.08	330 ± 18 ms	~ 660	~ 76M
Expert-VADv2	0.10	0.28	78.18	58.34	89.32	258 ± 22 ms	~ 573	~ 105M

Table 1: **Overall performance comparison.** ExpertAD achieves improved planning effectiveness and lower inference latency compared to baseline models. Performance is measured on a single NVIDIA GeForce RTX 3090 GPU.

Approach	Merge ↑	Overtake ↑	EmgBrake ↑	GiveWay ↑	Tsign ↑
UniAD	12.66	13.33	20.00	10.00	13.23
Expert-UniAD	27.38	23.67	51.67	20.00	40.93
VAD	8.89	20.44	18.64	20.00	18.66
Expert-VAD	22.25	26.38	49.33	20.00	51.53
VADv2	36.25	48.33	74.28	50.00	60.14
Expert-VADv2	40.44	48.33	78.42	40.00	65.78

Table 2: **Multi-skill capabilities in rare scenarios.** ExpertAD shows improved overall performance in Overtake, Merge, and T-sign scenarios, while maintaining comparable performance in EmgBrake and Giveaway scenarios.

Approach	τ	DS ↑	SR ↑	RC ↑
UniAD	-	44.62	14.09	68.68
UniAD + PA	32	40.35	15.00	58.94
UniAD + PA	64	48.25	16.93	70.47
UniAD + PA	128	52.53	18.41	76.73
UniAD + PA	256	43.31	16.50	64.83

Table 3: **PA results.** UniAD with PA achieves its best results with $\tau = 128$.

(computational complexity), and Params (model size) as efficiency metrics. Paired t-tests across benchmarks and metrics demonstrate that ExpertAD achieves statistically significant improvements over all baseline ADSs, with an average p-value of 0.026 ($p < 0.05$). All results are averaged over five independent runs.

Joint Results

As shown in Table 1, ExpertAD enhances both planning effectiveness and inference efficiency across all baselines. Compared to UniAD, our method achieves a 23% reduction in collision rates, a 14% decrease in L2 errors, and a 1.2× speedup. Expert-VAD and Expert-VADv2 exhibit similar improvements, reducing collision rates by 21% and 17%, lowering L2 errors by 9% and 15%, and achieving 1.4× and

Approach	Top-K	DS ↑	SR ↑	RC ↑
UniAD	-	44.62	14.09	68.68
UniAD + MoSE	Top-8	46.24	17.27	67.54
UniAD + MoSE	Top-4	49.32	18.41	72.03

Table 4: **MoSE results.** Top-4 experts contribute to the best planning performance over baseline UniAD.

1.3× speedups, respectively. In closed-loop evaluation, ExpertAD further improves DS, SR, and RC by 16%, 22%, and 14%, averaged over the three ADSs. ExpertAD incurs minimal parameter overhead while significantly lowering FLOPs, achieving a balanced trade-off between accuracy and efficiency. Unlike prior methods that favor speed at the cost of planning quality, ExpertAD offers joint gains in both.

Bench2Drive evaluates ADSs in five scenario groups. As shown in Table 2, ExpertAD excels in Emergency Braking situations (*e.g.*, PedestrianCrossing, ParkingCutIn), Merging situations (*e.g.*, HighwayCutIns, JunctionTurns), and Traffic Signs (*e.g.*, StopSigns, TrafficLights), where rich perceptual information from pedestrians, crossroads, and traffic lights provides critical context. However, for Overtaking scenarios (*e.g.*, Accident, ConstructionZones) and Giving Way scenarios (*e.g.*, YieldingtoEmergencyVehicles), ExpertAD shows only marginal or inconsistent improvements. These scenarios demand complex, human-like reasoning,

Approach	MLP()	ADD()	AMOTA \uparrow	AMOTP \downarrow
UniAD			0.388	1.304
Expert-UniAD	✓		0.384	1.306
	✓	✓	0.390	1.298
	✓	✓	0.404	1.277

Table 5: **Ablation results for PA components.** PA with an MLP layer introduces non-linearity to transformed features, while PA using addition function preserves feature stability.

Approach	Router	Sparse Attention	L2 (m) \downarrow				Col. Rate (%) \downarrow				Latency \downarrow
			1 s	2 s	3 s	Avg.	1 s	2 s	3 s	Avg.	
Expert-UniAD	✓		0.46	0.91	1.57	0.98	0.05	0.20	0.55	0.27	623 \pm 35ms
			0.34	0.76	1.47	0.86	0.03	0.16	0.45	0.21	591 \pm 30ms
	✓	✓	0.46	0.95	1.64	1.02	0.05	0.29	0.52	0.29	476 \pm 25ms
			0.42	0.85	1.41	0.89	0.07	0.22	0.43	0.24	445 \pm 20ms
Expert-VAD	✓		0.58	1.10	1.69	1.12	0.07	0.22	0.79	0.36	289 \pm 28ms
			0.53	1.05	1.67	1.08	0.06	0.16	0.73	0.32	256 \pm 22ms
	✓	✓	0.62	1.14	1.70	1.15	0.12	0.28	0.84	0.41	180 \pm 20ms
			0.53	1.07	1.69	1.10	0.07	0.22	0.74	0.34	157 \pm 23ms

Table 6: **Ablation results for MoSE components.** MoSE with Router routes the ego query to the top-4 experts, while MoSE without Router activates all experts. All experiments are conducted using PA. Router notably enhances planning metrics, while Sparse Attention significantly reduces latency.

suggesting that rule-based fallback systems are still needed.

Modular Results

Following the sequential order of the PA and MoSE module in ExpertAD, we evaluate the effectiveness of each module in comparison to the baseline models.

PA Results. We evaluate the hyperparameter τ , which controls the number of selected BEV features in the PA module. As shown in Table 3, performance improves as τ increases, peaking at $\tau = 128$ with gains of 17% (DS), 30% (SR), and 11% (RC) over the baseline. Larger τ values capture richer features, while overly large values (e.g., $\tau = 256$) may introduce redundancy.

MoSE Results. We compare three UniAD settings, baseline (no expert activation), Top-8 (all experts), and Top-4 (sparse activation). As shown in Table 4, Top-8 achieves modest DS gains, confirming the benefit of expert specialization. However, activating all experts leads to cautious decisions that reduce violations but slightly lower route completion. Top-4 further improves DS, SR, and RC by 11%, 30%, and 5%, respectively, as MoSE effectively reduces interference by selecting domain-relevant experts.

Ablation Study

We conduct ablation study on the PA and MoSE module respectively in Expert-UniAD and Expert-VAD. As VADv2 extends VAD with probabilistic modeling in planning, we exclude Expert-VADv2 from this study. To avoid the runtime overhead introduced by closed-loop evaluation, following experiments are performed under the open-loop setting.

PA Ablation. Since VAD does not implement a perception head, we compare only against UniAD for multi-object tracking. We adopt standard metrics (Hu et al. 2022), *i.e.*, Average Multi-Object Tracking Accuracy (AMOTA) and Average Multi-Object Tracking Precision (AMOTP) to assess the effectiveness of multi-object tracking. As shown in Table 5, removing the addition function reduces performance below the baseline, while adding the MLP layer notably improves AMOTA, highlighting its effectiveness in enhancing task-relevant features.

MoSE Ablation. Table 6 demonstrates that the Router effectively reduces L2 error and collision rate by mitigating task interference through selective expert activation. Sparse Attention further lowers latency with little or no loss in planning performance, and improves long-term prediction (e.g., 3s) by filtering out irrelevant tokens. Notably, Expert-UniAD with both Router and Sparse Attention achieves a significant latency reduction of 178ms, while improving planning effectiveness by 0.09m (L2 error) and 0.03% (collision rate). Expert-VAD shows similar benefits with a 132ms latency reduction and gains of 0.02m and 0.02%, confirming a favorable trade-off between efficiency and effectiveness.

Generalization Evaluation

To assess generalization, we conduct cross-city tests on nuScenes by training on one city (Boston / Singapore) and testing on the other. As shown in Table 7, Expert-UniAD reduces collision rates by 17% when trained on Boston and tested on Singapore. Conversely, when trained on Singapore, it achieves 13% lower L2 errors and 14% fewer collisions on Boston. Expert-VAD shows similar improvements across

Approach	Train on Boston		Train on Singapore	
	Avg. L2 ↓	Avg. Col. ↓	Avg. L2 ↓	Avg. Col. ↓
UniAD	1.26	0.29	1.24	0.66
Expert-UniAD	1.24	0.24	1.08	0.57
VAD	1.41	0.38	1.38	0.75
Expert-VAD	1.38	0.36	1.33	0.69

Table 7: **New-city generalization results.** We train the model on Boston and test it on Singapore, and vice versa, to evaluate the generalization capabilities.



Figure 3: **ExpertAD recovery visualization.** (a) UniAD misses a traffic officer on the right, leading the vehicle to drift toward them. (b) Expert-UniAD detects the officer, adjusts the trajectory, and completes the lane change safely.

both settings. While Boston poses more challenging road conditions (higher open-loop metrics), models trained on Singapore still achieve greater gains on Boston, highlighting the strong generalization of ExpertAD.

Qualitative Study

Fig. 3 illustrates the effectiveness of ExpertAD in a complex scenario. Here, UniAD fails to detect a traffic officer positioned ahead on the right, resulting in a planned path that veers toward the officer. In contrast, Expert-UniAD accurately detects the officer and immediately adjusts the route within the current frame, ensuring a safe trajectory.

Conclusions

We have proposed ExpertAD to improve the planning effectiveness and inference efficiency of ADSs with a Mixture of Experts architecture. Our evaluation has demonstrated the effectiveness and efficiency of ExpertAD, and the contribution of each module in ExpertAD. Further, ExpertAD also exhibits strong generalization capability. Qualitative results and case studies further highlight its capabilities.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No. 92582205).

References

- Bolya, D.; Fu, C.-Y.; Dai, X.; Zhang, P.; Feichtenhofer, C.; and Hoffman, J. 2022. Token merging: Your vit but faster. *arXiv preprint arXiv:2210.09461*.
- Caesar, H.; Bankiti, V.; Lang, A. H.; Vora, S.; Liong, V. E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; and Beijbom, O. 2019. nuScenes: A Multimodal Dataset for Autonomous Driving. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11618–11628.
- Casas, S.; Sadat, A.; and Urtasun, R. 2021. Mp3: A unified model to map, perceive, predict and plan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14403–14412.
- Chen, S.; Jiang, B.; Gao, H.; Liao, B.; Xu, Q.; Zhang, Q.; Huang, C.; Liu, W.; and Wang, X. 2024. Vadv2: End-to-end vectorized autonomous driving via probabilistic planning. *arXiv preprint arXiv:2402.13243*.
- Chen, X.; Huang, L.; Ma, T.; Fang, R.; Shi, S.; and Li, H. 2025. SOLVE: Synergy of Language-Vision and End-to-End Networks for Autonomous Driving. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 12068–12077.
- Daxberger, E.; Weers, F.; Zhang, B.; Gunter, T.; Pang, R.; Eichner, M.; Emmersberger, M.; Yang, Y.; Toshev, A.; and Du, X. 2023. Mobile V-MoEs: Scaling Down Vision Transformers via Sparse Mixture-of-Experts. *arXiv preprint arXiv:2309.04354*.
- Dosovitskiy, A.; Ros, G.; Codevilla, F.; Lopez, A.; and Koltun, V. 2017. CARLA: An Open Urban Driving Simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*, 1–16.
- Du, N.; Huang, Y.; Dai, A. M.; Tong, S.; Lepikhin, D.; Xu, Y.; Krikun, M.; Zhou, Y.; Yu, A. W.; Firat, O.; et al. 2022. Glam: Efficient scaling of language models with mixture-of-experts. In *International Conference on Machine Learning*, 5547–5569. PMLR.
- Fang, S.; and Choromanska, A. 2020. Multi-modal experts network for autonomous driving. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 6439–6445. IEEE.

- Fedus, W.; Zoph, B.; and Shazeer, N. 2022. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120): 1–39.
- Feng, K.; Li, C.; Ren, D.; Yuan, Y.; and Wang, G. 2024. On the Road to Portability: Compressing End-to-End Motion Planner for Autonomous Driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15099–15108.
- Feng, R.; Xi, N.; Chu, D.; Wang, R.; Deng, Z.; Wang, A.; Lu, L.; Wang, J.; and Huang, Y. 2025. Artemis: Autoregressive end-to-end trajectory planning with mixture of experts for autonomous driving. *arXiv preprint arXiv:2504.19580*.
- Gao, R.; Chen, K.; Xie, E.; Hong, L.; Li, Z.; Yeung, D.-Y.; and Xu, Q. 2023. Magicdrive: Street view generation with diverse 3d geometry control. *arXiv preprint arXiv:2310.02601*.
- Hu, S.; Chen, L.; Wu, P.; Li, H.; Yan, J.; and Tao, D. 2022. St-p3: End-to-end vision-based autonomous driving via spatial-temporal feature learning. In *European Conference on Computer Vision*, 533–549. Springer.
- Hu, Y.; Yang, J.; Chen, L.; Li, K.; Sima, C.; Zhu, X.; Chai, S.; Du, S.; Lin, T.; Wang, W.; et al. 2023. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17853–17862.
- Jia, X.; Gao, Y.; Chen, L.; Yan, J.; Liu, P. L.; and Li, H. 2023. Driveadapter: Breaking the coupling barrier of perception and planning in end-to-end autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7953–7963.
- Jia, X.; Yang, Z.; Li, Q.; Zhang, Z.; and Yan, J. 2024. Bench2Drive: Towards Multi-Ability Benchmarking of Closed-Loop End-To-End Autonomous Driving. In *Proceedings of the 38th Advances in Neural Information Processing Systems*.
- Jia, X.; You, J.; Zhang, Z.; and Yan, J. 2025. Drive-transformer: Unified transformer for scalable end-to-end autonomous driving. *arXiv preprint arXiv:2503.07656*.
- Jiang, A. Q.; Sablayrolles, A.; Roux, A.; Mensch, A.; Savary, B.; Bamford, C.; Chaplot, D. S.; Casas, D. d. l.; Hanna, E. B.; Bressand, F.; et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Jiang, B.; Chen, S.; Xu, Q.; Liao, B.; Chen, J.; Zhou, H.; Zhang, Q.; Liu, W.; Huang, C.; and Wang, X. 2023. Vad: Vectorized scene representation for efficient autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8340–8350.
- Kim, D.; Woo, S.; Lee, J.-Y.; and Kweon, I. S. 2020. Video panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9859–9868.
- Kirillov, A.; He, K.; Girshick, R.; Rother, C.; and Dollár, P. 2019. Panoptic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9404–9413.
- Lei, T.; Bai, J.; Brahma, S.; Ainslie, J.; Lee, K.; Zhou, Y.; Du, N.; Zhao, V.; Wu, Y.; Li, B.; et al. 2023. Conditional adapters: Parameter-efficient transfer learning with fast inference. *Advances in Neural Information Processing Systems*, 36: 8152–8172.
- Li, P.; and Jin, J. 2022. Time3d: End-to-end joint monocular 3d object detection and tracking for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3885–3894.
- Li, Z.; Wang, S.; Lan, S.; Yu, Z.; Wu, Z.; and Alvarez, J. M. 2025. Hydra-next: Robust closed-loop driving with open-loop training. *arXiv preprint arXiv:2503.12030*.
- Li, Z.; Wang, W.; Li, H.; Xie, E.; Sima, C.; Lu, T.; Qiao, Y.; and Dai, J. 2022a. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *European conference on computer vision*, 1–18. Springer.
- Li, Z.; Wang, W.; Xie, E.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; Luo, P.; and Lu, T. 2022b. Panoptic segformer: Delving deeper into panoptic segmentation with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1280–1289.
- Liang, T.; Xie, H.; Yu, K.; Xia, Z.; Lin, Z.; Wang, Y.; Tang, T.; Wang, B.; and Tang, Z. 2022. Bevfusion: A simple and robust lidar-camera fusion framework. *Advances in Neural Information Processing Systems*, 35: 10421–10434.
- Morra, L.; Biondo, A.; Poerio, N.; and Lamberti, F. 2023. MIXO: Mixture of experts-based visual odometry for multicamera autonomous systems. *IEEE Transactions on Consumer Electronics*, 69(3): 261–270.
- Nazeri, M. H.; and Bohlouli, M. 2021. Exploring reflective limitation of behavior cloning in autonomous vehicles. In *2021 IEEE International Conference on Data Mining (ICDM)*, 1252–1257. IEEE.
- Ohn-Bar, E.; Prakash, A.; Behl, A.; Chitta, K.; and Geiger, A. 2020. Learning situational driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11296–11305.
- Pini, S.; Perone, C. S.; Ahuja, A.; Ferreira, A. S. R.; Nien-dorf, M.; and Zagoruyko, S. 2023. Safe real-world autonomous driving by learning to predict and plan with a mixture of experts. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 10069–10075. IEEE.
- Pomerleau, D. A. 1988. Alvin: An autonomous land vehicle in a neural network. *Advances in neural information processing systems*, 1.
- Prakash, A.; Chitta, K.; and Geiger, A. 2021. Multi-modal fusion transformer for end-to-end autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7077–7087.
- Rao, Y.; Zhao, W.; Liu, B.; Lu, J.; Zhou, J.; and Hsieh, C.-J. 2021. Dynamicvit: Efficient vision transformers with dynamic token sparsification. *Advances in neural information processing systems*, 34: 13937–13949.
- Riquelme, C.; Puigcerver, J.; Mustafa, B.; Neumann, M.; Jenatton, R.; Susano Pinto, A.; Keysers, D.; and Houlsby,

- N. 2021. Scaling vision with sparse mixture of experts. *Advances in Neural Information Processing Systems*, 34: 8583–8595.
- Shao, H.; Hu, Y.; Wang, L.; Song, G.; Waslander, S. L.; Liu, Y.; and Li, H. 2024. Lmdrive: Closed-loop end-to-end driving with large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15120–15130.
- Shao, H.; Wang, L.; Chen, R.; Waslander, S. L.; Li, H.; and Liu, Y. 2023. Reasonnet: End-to-end driving with temporal and global reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13723–13733.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Sima, C.; Renz, K.; Chitta, K.; Chen, L.; Zhang, H.; Xie, C.; Luo, P.; Geiger, A.; and Li, H. 2023. Drivelm: Driving with graph visual question answering. *arXiv preprint arXiv:2312.14150*.
- Sun, Y.; Wang, X.; Zhang, Y.; Tang, J.; Tang, X.; and Yao, J. 2023. Interpretable End-to-End Driving Model for Implicit Scene Understanding. In *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, 2874–2880. IEEE.
- Team, L.-M. 2023. Llama-moe: Building mixture-of-experts from llama with continual pre-training.
- Tian, X.; Gu, J.; Li, B.; Liu, Y.; Wang, Y.; Zhao, Z.; Zhan, K.; Jia, P.; Lang, X.; and Zhao, H. 2024. Drivevlm: The convergence of autonomous driving and large vision-language models. *arXiv preprint arXiv:2402.12289*.
- Toromanoff, M.; Wirbel, E.; and Moutarde, F. 2020. End-to-end model-free reinforcement learning for urban driving using implicit affordances. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7153–7162.
- Wang, D.; Lu, Y.; Chen, B.; Hao, S.; Jiang, H.; Tian, Y.; and Peng, X. 2025a. Argus: Resilience-Oriented Safety Assurance Framework for End-to-End ADSs. *arXiv preprint arXiv:2511.09032*.
- Wang, H.; Cai, P.; Sun, Y.; Wang, L.; and Liu, M. 2021. Learning interpretable end-to-end vision-based motion planning for autonomous driving with optical flow distillation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 13731–13737. IEEE.
- Wang, T.; Zhang, C.; Qu, X.; Li, K.; Liu, W.; and Huang, C. 2025b. DiffAD: A Unified Diffusion Modeling Approach for Autonomous Driving. *arXiv preprint arXiv:2503.12170*.
- Wang, X.; Zhu, Z.; Huang, G.; Chen, X.; Zhu, J.; and Lu, J. 2023. Drivedreamer: Towards real-world-driven world models for autonomous driving. *arXiv preprint arXiv:2309.09777*.
- Wang, Y.; He, J.; Fan, L.; Li, H.; Chen, Y.; and Zhang, Z. 2024. Driving into the future: Multiview visual forecasting and planning with world model for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14749–14759.
- Xiao, Y.; Codevilla, F.; Bustamante, D. P.; and Lopez, A. M. 2023. Scaling self-supervised end-to-end driving with multi-view attention learning. *arXiv preprint arXiv:2302.03198*, 2.
- Xu, Z.; Zhang, Y.; Xie, E.; Zhao, Z.; Guo, Y.; Wong, K.-Y. K.; Li, Z.; and Zhao, H. 2024. Drivept4: Interpretable end-to-end autonomous driving via large language model. *IEEE Robotics and Automation Letters*.
- Yang, K.; Ma, E.; Peng, J.; Guo, Q.; Lin, D.; and Yu, K. 2023. Bevcontrol: Accurately controlling street-view elements with multi-perspective consistency via bev sketch layout. *arXiv preprint arXiv:2308.01661*.
- Ye, T.; Jing, W.; Hu, C.; Huang, S.; Gao, L.; Li, F.; Wang, J.; Guo, K.; Xiao, W.; Mao, W.; et al. 2023. Fusionad: Multi-modality fusion for prediction and planning tasks of autonomous driving. *arXiv preprint arXiv:2308.01006*.
- Yin, H.; Vahdat, A.; Alvarez, J. M.; Mallya, A.; Kautz, J.; and Molchanov, P. 2022. A-vit: Adaptive tokens for efficient vision transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10809–10818.
- Yuan, R.; Abdel-Aty, M.; Xiang, Q.; Wang, Z.; and Gu, X. 2023. A temporal multi-gate mixture-of-experts approach for vehicle trajectory and driving intention prediction. *IEEE Transactions on Intelligent Vehicles*.
- Zeng, F.; Dong, B.; Zhang, Y.; Wang, T.; Zhang, X.; and Wei, Y. 2022. Motr: End-to-end multiple-object tracking with transformer. In *European Conference on Computer Vision*, 659–675. Springer.
- Zhang, E.; Dai, X.; Lv, Y.; and Miao, Q. 2024a. MiniDrive: More Efficient Vision-Language Models with Multi-Level 2D Features as Text Tokens for Autonomous Driving. *arXiv preprint arXiv:2409.07267*.
- Zhang, T.; Chen, X.; Wang, Y.; Wang, Y.; and Zhao, H. 2022. Mutr3d: A multi-camera tracking framework via 3d-to-2d queries. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4537–4546.
- Zhang, Y.; Gong, S.; Xiong, K.; Ye, X.; Tan, X.; Wang, F.; Huang, J.; Wu, H.; and Wang, H. 2024b. BEVWorld: A Multimodal World Model for Autonomous Driving via Unified BEV Latent Space. *arXiv preprint arXiv:2407.05679*.
- Zhang, Z.; Liniger, A.; Dai, D.; Yu, F.; and Van Gool, L. 2021. End-to-end urban driving by imitating a reinforcement learning coach. In *Proceedings of the IEEE/CVF international conference on computer vision*, 15222–15232.