

# ARTEMIS: Autoregressive End-to-End Trajectory Planning with Mixture of Experts for Autonomous Driving

Renju Feng<sup>†</sup>, Ning Xi<sup>†</sup>, Duanfeng Chu\*, Rukang Wang, Zejian Deng, Anzheng Wang, Liping Lu, Jinxiang Wang, Yanjun Huang

**Abstract**—This paper presents ARTEMIS, an end-to-end autonomous driving framework that combines autoregressive trajectory planning with Mixture-of-Experts (MoE). Traditional modular methods suffer from error propagation, while existing end-to-end models typically employ static one-shot inference paradigms that inadequately capture the dynamic changes of the environment. ARTEMIS takes a different method by generating trajectory waypoints sequentially, preserves critical temporal dependencies while dynamically routing scene-specific queries to specialized expert networks. It effectively relieves trajectory quality degradation issues encountered when guidance information is ambiguous, and overcomes the inherent representational limitations of singular network architectures when processing diverse driving scenarios. Additionally, we use a lightweight batch reallocation strategy that significantly improves the training speed of the Mixture-of-Experts model. Through experiments on the NAVSIM dataset, ARTEMIS exhibits superior competitive performance, achieving 86.9 PDMS and 83.1 EPDMS with ResNet-34 backbone, demonstrates state-of-the-art performance on multiple metrics. Code will be available under <https://github.com/Lg0914/ARTEMIS>.

**Index terms**— Autonomous Vehicle Navigation; Integrated Planning and Learning; Deep Learning Methods.

## I. INTRODUCTION

**A**UTONOMOUS driving has experienced rapid development over the past few decades. Traditional modular methods compartmentalize autonomous driving tasks into discrete modules such as perception, prediction, and planning

Manuscript received: May 1, 2025; revised September 23, 2025; accepted November 6, 2025. This paper was recommended for publication by Editor T. Asfour and A. Banerjee upon evaluation of the reviewers’ comments.

This work is supported in part by the National Natural Science Foundation of China (52472438), the Key R&D Program of Hubei Province (2024BAB033).

Renju Feng, Ning Xi, Duanfeng Chu, Rukang Wang, and Anzheng Wang are with the Intelligent Transportation Systems Research Center, Wuhan University of Technology, Wuhan 430063, China (e-mail: fengrenju; xining095; chudf; wangrk; wanganzheng@whut.edu.cn).

Zejian Deng is with the Department of Data and Systems Engineering, The University of Hong Kong, Hong Kong SAR, 999077, China (e-mail: z49deng@hku.hk).

Liping Lu is with the School of Computer Science and Artificial Intelligence, Wuhan University of Technology, Wuhan 430070, China (e-mail: luliping@whut.edu.cn).

Jinxiang Wang is with the School of Mechanical Engineering, Southeast University, Nanjing 211189, China (e-mail: wangjx@seu.edu.cn).

Yanjun Huang is with the School of Automotive Studies, Tongji University, Shanghai 201804, China (e-mail: yanjun\_huang@tongji.edu.cn).

<sup>†</sup>Equal contribution.

\*Corresponding author.

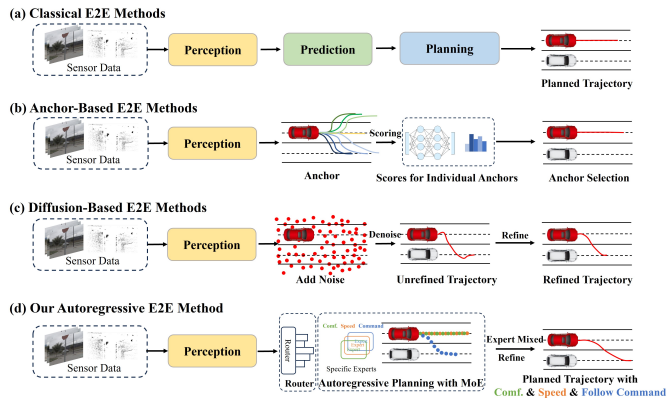


Fig. 1. Four different end-to-end architectures. (a) Classical end-to-end methods [4], [5], [12], [13]. (b) Anchor-based end-to-end methods [14]–[16]. (c) Diffusion-based end-to-end methods [17]. (d) Our proposed method.

[1]–[3]. However, cumulative errors and complex interdependencies between these modules can become constrained at pre-defined interfaces. End-to-end models overcome these issues by mapping raw sensor data directly to planned trajectories or control signals [4]–[6]. However, their static, one-shot inference paradigm often fails to capture the dynamic evolution of the environment (Fig.1 (a)). In contrast, autoregressive methods generate trajectories sequentially, preserving temporal coherence and allowing for adaptive decision-making based on previously planned segments. Autoregressive models have been widely applied in the field of trajectory prediction [7]–[9]. Recent research has begun to explore unified frameworks that simultaneously accomplish world model construction and complete or partial trajectory planning tasks through autoregressive modeling methods [10], [11]. However, the current problem is that single-network end-to-end models still struggle to adequately capture and adapt to diverse driving scenarios.

To tackle the inherent complexity in trajectory planning, researchers are increasingly adopting sophisticated architectures like the Mixture of Experts (MoE) framework [18]. MoE employs multiple specialized expert networks along with intelligent routing mechanisms to dynamically distribute and process inputs—a strategy that has yielded significant success in large-scale language models [19], [20]. In autonomous driving, the planning trajectories generated by end-to-end models inherently encompass multiple potential behavioral modalities, reflecting the fundamental uncertainty in driving behavior. Drivers may select from several reasonable future

actions under identical environmental conditions (Fig.1 (b)) [14]–[16], and traditional single-network architectures struggle to accurately characterize this intrinsic diversity of behavioral patterns. In contrast, MoE enables expert modules to focus on specific driving scenarios or behavioral patterns, thereby learning characteristic distributions of driving behavior without reliance on predefined guidance signals. This endogenous multimodal modeling method effectively circumvents potential trajectory quality degradation issues that may arise when guidance information deviates from actual conditions. Recently, diffusion-based models (Fig.1 (c)) [17], have introduced novel generative modeling paradigms to autonomous driving, demonstrating trajectory diversity. Although these methodologies have established state-of-the-art performance in end-to-end autonomous driving, they typically employ static paradigms that generate all trajectory points simultaneously (or through multiple denoising iterations), limiting their capacity to accurately capture the dynamic evolutionary characteristics of trajectory development. In comparison, autoregressive methods utilizing MoE demonstrate superior temporal sequence capturing capabilities, environmental adaptability, and practical utility, as they can operate without requiring strong prior constraints.

To address these issues, we propose ARTEMIS, Autoregressive End-to-End Trajectory Planning with Mixture of Experts for Autonomous Driving, as shown in Fig.2. ARTEMIS comprises three primary components: a perception module, an autoregressive planning module with MoE, and a trajectory refinement module. The perception module adopts the Transfuser [13], utilizing separate backbone networks for image and LiDAR data, ultimately fusing them into BEV feature representations. The autoregressive planning module with MoE progressively generates trajectory waypoints through sequential decision processes while dynamically selecting the most appropriate expert network for the current driving scenario. Additionally, it implements batch reallocation based on expert activation patterns. Finally, the trajectory refinement module processes and refines the autoregressive trajectory outputs. We conducted comprehensive evaluations of ARTEMIS using the NAVSIM dataset [21]. Our contributions can be summarized as follows:

(1) To the best of our knowledge, this study represents the first investigation to incorporate the Mixture-of-Experts (MoE) into end-to-end autonomous driving, effectively relieving trajectory quality degradation issues encountered in traditional methods when guidance information is ambiguous, as well as the inherent representational limitations of singular network architectures when processing diverse driving scenarios, through dynamic routing mechanisms and specialized expert network division.

(2) We propose an autoregressive end-to-end planning methodology that constructs trajectories sequentially through iterative decision processes, enabling precise modeling of strong temporal dependencies between trajectory waypoints.

(3) Our method achieves significant results on the large-scale real-world NAVSIM dataset. Using an identical ResNet-34 backbone, our method attains 86.9 PDMS under standard metrics and 83.1 EPDMS under extended evaluation metrics.

## II. METHOD

### A. Model Architecture

1) *Perception Module*: The perception module adheres to the design of the Transfuser [13] and is responsible for extracting features from raw sensor data. This module employs a multimodal fusion strategy to process both image and point cloud data simultaneously, thereby constructing a unified representation of the environment.

The module utilizes point cloud data  $D_l \in \mathbb{R}^{256 \times 256}$  and front-view image data  $D_I \in \mathbb{R}^{1024 \times 256 \times 3}$ , from which visual features (e.g.,  $F_{cam}$  and  $F_{lidar}$ ) are extracted using a series of convolution layers and a ResNet-34 backbone. Finally, a multimodal fusion mechanism is employed to integrate these features into a Bird’s-Eye View (BEV) feature representation  $F_{bev} \in \mathbb{R}^{C_{bev} \times 256}$ , where  $C_{bev}$  represents the number of BEV channels.

2) *Autoregressive Planning Module with Mixture-of-Experts*: Unlike traditional one-shot methods, this study employs an autoregressive strategy to incrementally build trajectories while integrating a Mixture-of-Experts (MoE) architecture. This design leverages both previous trajectory information and dynamically selected specialized expert networks based on scene characteristics.

Our analysis of the navtrain split reveals a significant imbalance in the distribution of driving commands (with over 20,000 samples for left-turn commands, fewer than 10,000 samples for right-turn commands, and more than 50,000 samples for forward driving commands). **Furthermore, we observed that a small proportion of the training samples exhibit discrepancies between the driving commands and the expert trajectories** (as illustrated in Fig.3). Relying solely on driving commands for expert selection may result in insufficient training data for certain experts and fail to capture the diversity of driving strategies. Ultimately, we adopt an endogenous routing multimodal modeling method to effectively relieve these challenges.

To avoid the model exhibiting a bias toward directly learning planning from historical ego trajectory, which may lead to causal confusion and other issues [22], [23]. We exclusively encode the current ego state  $s_0 \in \mathbb{R}^8$  (including control commands, 2D velocity, and acceleration) into a feature space  $Q_s$  via an MLP.

**Positional and Temporal Embedding.** It is essential to inform the model of the specific time step for which the trajectory point is being planned. For a future planning time step  $t$ , we employ an embedding layer to obtain the planning time embedding, denoted as  $TE_t$ . Similarly, we utilize a positional embedding to incorporate positional information into the planning sequence, represented by  $PE_t$ . It is important to note that we only add the positional embedding to the initialized planning sequence during the first autoregressive step, thus avoiding the accumulation of extraneous noise that may result from repeatedly adding the embedding information.

**Autoregressive Generation.** We first add positional embeddings to the complete current planning sequence for initialization. This sequence is then fed into a Transformer encoder with a padding mask  $M_t$  to update the planning queries  $Q_{1:t}$ .

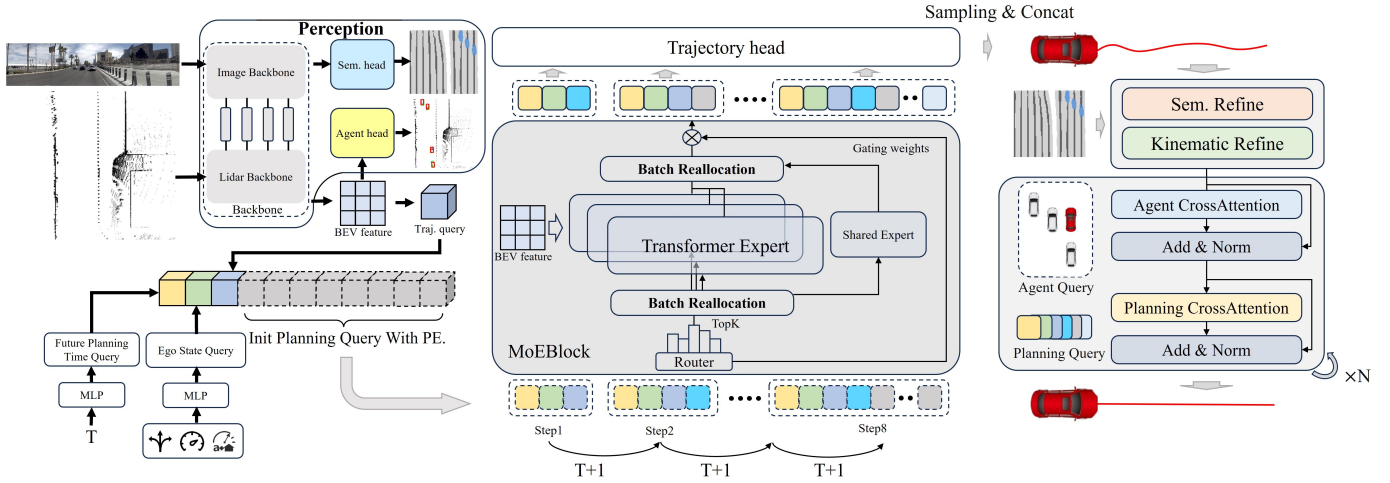


Fig. 2. The figure illustrates our method: The perception module utilizes separate backbone networks for image and LiDAR data, ultimately fusing them into BEV feature representations. The autoregressive planning module with MoE progressively generates trajectory waypoints through sequential decision processes while dynamically selecting the most appropriate expert network for the current environment. Additionally, it implements batch reallocation based on expert activation patterns. Finally, the trajectory refinement module processes and refines the autoregressive trajectory outputs.

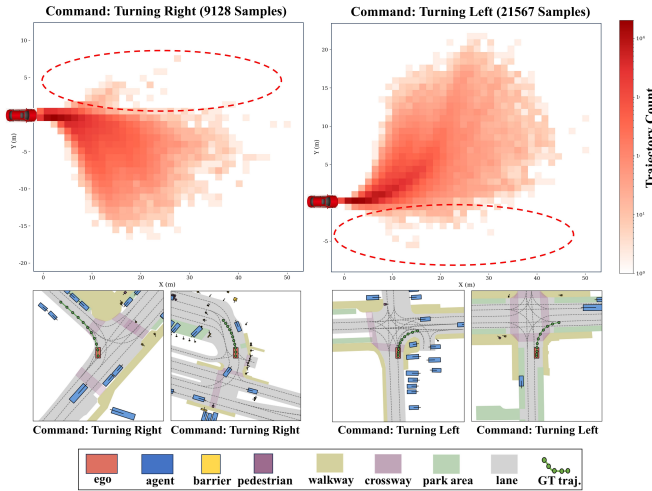


Fig. 3. The lower part of this figure shows several typical scenes from the navtrain split where control commands are inconsistent with the actual driving trajectories. The upper part visualizes the distribution of expert trajectories under right and left turn control commands through grid heatmaps, clearly showing trajectory deviations where some grids are inappropriately activated due to incorrect control commands.

The padding mask ensures that the planning query at the current time step interacts only with the historical planning queries in the sequence. By constructing a concatenated query that integrates the planning time information, ego state, and the current as well as historical planning queries, we have  $C_t = \text{Concat}(TE_t, Q_s, Q_{(1:t)})$ . This concatenated query is then fed into the MoEBlock with batch reallocation to interact with the BEV feature and obtain the planning query for the current time step.

$$Q_{t+1} = \text{MoEBlock}(F_{bev}, C_t) \quad (1)$$

Finally, to further characterize the inherent uncertainty in driving behavior, we adopt a probabilistic modeling method. A multi-layer MLP network is employed to predict the distribution of multimodal trajectory points, encompassing both position and heading. Ultimately, a trajectory point at the

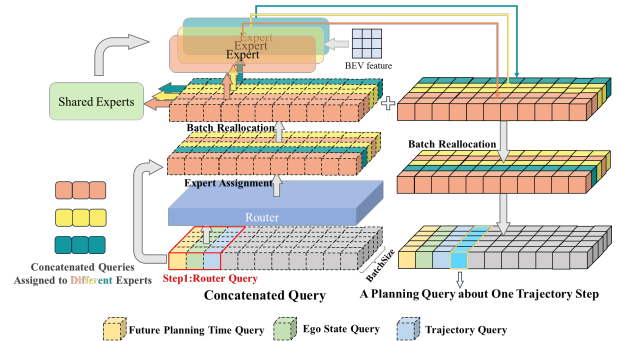


Fig. 4. **Batch Reallocation MoE**. This module intelligently allocates expert resources to concatenated queries from different samples through the routing network, and employs batch reorganization to group sample data that activate the same experts together, thereby improving computational efficiency. The system subsequently uses shared expert networks and different specialized expert networks to process and fuse BEV feature, then restores the original sample order through inverse batch reorganization after processing. Finally, the module extracts the query corresponding to the next position from the processed concatenated query, which serves as the planning query input for the next trajectory point supplied to the trajectory planning head.

current planning time step is generated by sampling from the predicted distribution  $y_t \sim \mathcal{N}(\mu_t, \sigma_t^2)$ . After the trajectory point is generated, the latest planning query corresponding to the current time step is updated in the planning query sequence. Finally, the individual trajectory points are concatenated to form the initial trajectory.

**Batch Reallocation MoE**. This module, as shown in Fig.4, includes both shared experts  $E_{\text{shared}}$  and domain-specific experts  $E_{\text{private}}$ . We use an efficient batch reallocation strategy that significantly enhances computational efficiency, particularly when processing large-scale data. To ensure a fixed query dimensionality for the routing network, we remove the historical planning queries from the concatenated query, forming the routing query  $Q_r = \text{Concat}(TE_t, Q_s, Q_t)$ . The router network  $\mathcal{R}$  consists of two MLPs and is used to compute the expert assignment scores  $g_t$ . The first MLP performs dimensionality reduction on input features, while the second maps compressed features to expert scores. This two-stage

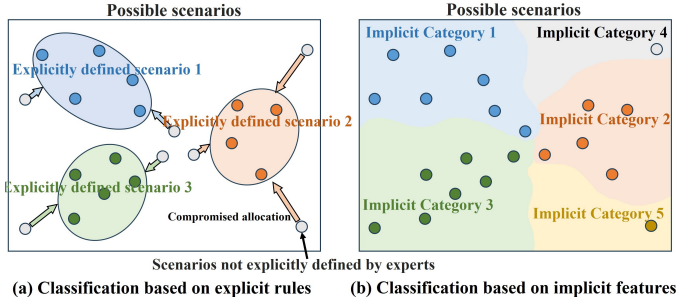


Fig. 5. **A conceptual comparison between explicit and implicit routing strategies.** (a) Explicit routing, based on human-defined classification rules, forces scenarios that fall outside these rules (white dots) into ill-fitting categories, resulting in “compromised allocations.” (b) The implicit routing allows the model to learn a data-driven partitioning of the entire feature space, enabling it to handle all cases with greater flexibility and robustness.

design efficiently balances computational cost with routing decision quality.

To enable the model to focus on the most relevant experts during training, we also adopt a sparse activation strategy by selecting only the top  $k$  experts with the highest score. For each selected expert index sequence vector  $a_i$ , we perform a series of operations including batch sorting, data reorganization, and block identification. Specifically, we first sort the batch samples according to the expert indices, resulting in a sorting function  $\pi_i$ . Subsequently, the BEV feature and the concatenated queries are restructured based on this sorting function as follows:

$$F_{\text{bev}}^{(i)} = \pi_i(F_{\text{bev}}) \quad \text{and} \quad Q_t^{(i)} = \pi_i(Q_t). \quad (2)$$

Based on expert indexing patterns, consecutive blocks with identical expert indices are identified and matched:

$$\{E_j^i, n_j^i\}_{j=1}^J = UC(\pi_i) \quad (3)$$

where  $J$  represents the total number of blocks,  $E_j^i$  denotes the consecutive block corresponding to the  $j_{th}$  expert in the  $i_{th}$  expert index sequence vector, and  $n_j^i$  represents the corresponding block size as determined by the unique consecutive function. The appropriate expert network is applied to each block, where  $F_{\text{bev}}^{(i,j)}$  and  $C_t^{(i,j)}$  constitute the reorganized expert input data for block  $E_j^i$  under the expert index sequence vector  $a_i$ .

$$F_{\text{bev}}^{(i,j)} = F_{\text{bev}}^{(i)}[b_j^i:(b_j^i + n_j^i)] \quad (4)$$

$$C_t^{(i,j)} = C_t^{(i)}[b_j^i:(b_j^i + n_j^i)] \quad (5)$$

$$O_{i,j} = E_{\text{private}}^j(F_{\text{bev}}^{(i,j)}, C_t^{(i,j)}) + E_{\text{shared}}(F_{\text{bev}}^{(i,j)}, C_t^{(i,j)}) + C_t^{(i,j)} \quad (6)$$

where  $b_j^i$  denotes the starting position of the consecutive block corresponding to the  $j_{th}$  expert in the  $i_{th}$  expert index sequence vector.

Finally, the expert processing results are restored to their original batch ordering, and outputs from different experts are fused according to their corresponding weights  $g_{ij}$  to derive the final composite output  $O_f$ .

$$\pi_i(\hat{O}_{i,j}) = O_{i,j} \quad \text{and} \quad O_f = \sum_i \sum_j (g_{ij} \times \hat{O}_{i,j}) \quad (7)$$

**Data-Driven Implicit Routing.** Our MoE architecture employs an implicit routing mechanism, deliberately avoiding an explicit, rule-based approach, as shown in Fig.5. While explicit routing using scene labels offers interpretability, it is often brittle. Such rule-based assignment methods struggle to cover the full spectrum of complex and long-tail driving scenarios, potentially leading to suboptimal expert allocation in undefined situations. We therefore used data-driven implicit router by removing hard constraints and allowing it to learn an optimal, fine-grained partitioning of the feature space directly from high-dimensional scene features. This approach prioritizes the model’s autonomous learning capabilities and adaptability over reliance on potentially incomplete, human-defined classification rules.

3) *Trajectory Refinement Module:* Considering that driving scenarios are inherently complex and often contain information that is unrelated to planning or even constitutes noise, we introduce a trajectory refinement module. The refinement process ensures that the final trajectory  $Y$  satisfies kinematic constraints, avoids obstacles, and maintains smoothness. This process is divided into two stages: semantic kinematic optimization and cross-attention refinement.

**Semantic Kinematic Optimization.** In the semantic optimization stage, we extract features from the BEV semantic map  $S_{\text{bev}}$  to obtain semantic features  $F_{\text{sem}} = \phi_{\text{sem}}(S_{\text{bev}})$ , where  $\phi_{\text{sem}}$  is a semantic encoder that is primarily implemented using a multi-layer convolutional network. We encode the initial trajectory using a GRU network to obtain the trajectory features  $h_{\text{traj}}$ , and subsequently feed both the semantic features  $F_{\text{sem}}$  and  $h_{\text{traj}}$  into an optimizer network  $\phi_{\text{optimizer}}$ , yielding the joint representation  $F_{\text{combined}}$ .

The decoding part optimizes the trajectory point by point by fusing the features. For each trajectory point, a GRU and an output layer  $\phi_{\text{output}}$  are used to generate the optimized point  $y'_j$ . In the kinematic optimization stage, multiple explicit constraints are employed to optimize the generated trajectory points, including both smoothness constraints and kinematic constraints, with learnable weights assigned to these constraints. The final optimized points obtained from semantic kinematic optimization are denoted as  $\hat{y}_i$ .

**Cross-Attention Refinement.** We utilize a cascaded cross-attention module to further enhance the interaction between the trajectory and the scene context, enabling the trajectory features to interact with both the agent features and the ego planning features.

$$Y = \Psi_{\text{refine}}(\hat{y}_i, Q_{\text{agent}}, Q_{\text{ego}}) \quad (8)$$

where  $Q_{\text{agent}}$  represents the agent query features,  $Q_{\text{ego}}$  denotes the ego planning query features, and  $\Psi_{\text{refine}}$  is the refinement function.

4) *Training Loss:* Consistent with some end-to-end methods, we adopt a staged training method to mitigate training instability [4], [5], [14]. Specifically, the perception network and its auxiliary tasks, including semantic mapping and object detection, are first trained. Subsequently, the entire network is trained in an end-to-end method. This training strategy significantly enhances both the stability and the overall performance of the model. Furthermore, we opted not to employ the expert

balance loss commonly utilized in MoE architecture, as its application in datasets with imbalanced feature distributions could potentially impede the acquisition of specialized strategic knowledge by individual experts.

**Perception Stage Loss.** In the first stage, we focus on optimizing the auxiliary tasks related to perception. The total loss is defined as:

$$\mathcal{L}_{\text{perception}} = \lambda_{\text{sem}}\mathcal{L}_{\text{sem}} + \lambda_{\text{class}}\mathcal{L}_{\text{class}} + \lambda_{\text{box}}\mathcal{L}_{\text{box}} \quad (9)$$

where  $\mathcal{L}_{\text{sem}}$  is the cross-entropy loss for the BEV semantic map, and  $\mathcal{L}_{\text{class}}$  and  $\mathcal{L}_{\text{box}}$  are the agent classification and localization losses computed using the Hungarian matching algorithm.

**Negative Log-Likelihood Loss.** NLL loss and variance regularization loss are designed to minimize the deviation between the predicted trajectory and the ground truth.

$$\mathcal{L}_{\text{NLL}} = \frac{1}{H} \sum_{t=1}^H (\mathcal{L}_{\text{NLL}_t} + \lambda_{\text{var}} \cdot \mathcal{L}_{\text{var\_reg}_t}) \quad (10)$$

$$\mathcal{L}_{\text{NLL}_t} = \frac{1}{2} \sum_{d \in \{x, y, \theta\}} \left( \log(2\pi\sigma_{t,d}^2) + \frac{(y_{t,d}^{\text{gt}} - \mu_{t,d})^2}{\sigma_{t,d}^2} \right) \quad (11)$$

$$\mathcal{L}_{\text{var\_reg}_t} = \sum_{d \in \{x, y, \theta\}} (\log(\sigma_{t,d}^2))^2 \quad (12)$$

where  $H$  is the total number of prediction steps,  $t$  is the current autoregressive time step, and  $d$  iterates over the three dimensions of a waypoint  $(x, y, \theta)$ . The terms  $\mu_{t,d}$  and  $\sigma_{t,d}^2$  are the mean and variance predicted by the model for dimension  $d$  at time step  $t$ , while  $y_{t,d}^{\text{gt}}$  is the corresponding ground-truth value. Finally,  $\lambda_{\text{var}}$  is a hyperparameter.

**End-to-End Training Loss.** In this stage, the entire network is trained in end-to-end method. The overall loss is defined as:

$$\mathcal{L}_{\text{planning}} = \lambda_{\text{traj}}\mathcal{L}_{\text{traj}} + \lambda_{\text{NLL}}\mathcal{L}_{\text{NLL}} + \mathcal{L}_{\text{perception}} \quad (13)$$

where  $\mathcal{L}_{\text{traj}}$  denotes the planning  $L_1$  loss, and  $\mathcal{L}_{\text{NLL}}$  represents the negative log-likelihood loss.

### III. EXPERIMENTS

#### A. Dataset

We train and test our model on the NAVSIM dataset [21]. The training set comprises 1192 scenarios, while the test set contains 136 scenarios. Each sample in the dataset includes camera images from 8 viewpoints, LiDAR data fused from 5 sensors, map annotations, and 3D object bounding boxes, among other data. In the NAVSIM dataset, the model is required to use 4 frames of 2 seconds of historical and current data to plan a 4-second trajectory consisting of 8 future frames.

#### B. Evaluation Metrics

Numerous studies have demonstrated that performing a simple open-loop evaluation is insufficient to fully assess a model's performance [22], [24], while closed-loop evaluation are hindered by high computational costs and discrepancies between the simulator and the real world. NAVSIM provides an intermediary solution between these evaluation paradigms by introducing the Predictive Driving Model Score (PDMS) [21], which shows a high correlation with closed-loop metrics. The PDMS is computed based on five indicators: No-Collision

(NC), Drivable Area Compliance (DAC), Time-to-Collision (TTC), Comfort (C), and Ego Vehicle Progress (EP).

In addition to PDMS, NAVSIM also provides an extended benchmark, the Extended Predictive Driving Model Score (EPDMS) [16]. The extended score introduces two new weighted metrics (Lane Keeping, LK, and Extended Comfort, EC), two new multiplicative metrics (Driving Direction Compliance, DDC, and Traffic Light Compliance, TLC), along with a false alarm penalty filtering process.

#### C. Implementation Details

We employ the Transfuser [13] as our perception network, with ResNet34 as the backbone for feature extraction. In the autoregressive planning module that incorporates MoE, we configure  $E_{\text{private}} = 5$  domain-specific experts along with  $E_{\text{shared}} = 1$  shared expert. During forward propagation, the top  $K = 2$  experts with the highest score are selected for activation. The model is trained on the navtrain split using two A100 GPUs with a batch size of 128. The initial learning rate is set to  $2 \times 10^{-4}$  and the weight decay is set to  $1 \times 10^{-4}$ . The model performs 8 autoregressive steps, with each step outputting one trajectory point comprising the  $x$ ,  $y$  and heading. After all autoregressive steps, the generated trajectory is refined to form a 4-second planning trajectory at 2 Hz. During the training stage of the perception module, the loss weight coefficients  $\lambda_{\text{sem}}$ ,  $\lambda_{\text{class}}$ , and  $\lambda_{\text{box}}$  are set to 10, 10, and 5, respectively. In the end-to-end training stage, the weight coefficients  $\lambda_{\text{sem}}$ ,  $\lambda_{\text{class}}$ ,  $\lambda_{\text{box}}$ ,  $\lambda_{\text{traj}}$ ,  $\lambda_{\text{var}}$ , and  $\lambda_{\text{NLL}}$  are set to 2, 2, 1, 15, 0.01, and 0.5, respectively.

#### D. Main Results

**Quantitative Results.** On the Navtest benchmark, we benchmarked ARTEMIS against several state-of-the-art methods, which is summarized in Tab. I. With a ResNet-34 backbone, ARTEMIS attains 86.9 PDMS on the navtest split. Notably, ARTEMIS achieves a markedly superior EP, NC and C score, evidencing its strong trajectory planning capability and environmental adaptability. We further evaluated ARTEMIS on the Navtest benchmark with extended metrics (Tab. II). It demonstrates that our model significantly outperforms all other baselines using the same ResNet-34 backbone, achieving state-of-the-art results. Notably, our method substantially exceeds comparative methodologies on critical metrics such as TTC and EP. These results underscore ARTEMIS's robustness and superior performance across diverse evaluation criteria.

**Qualitative Results.** Fig. 6 demonstrates four representative driving scenarios sampled from the navtest dataset to qualitatively evaluate the proposed method. To highlight how individual experts respond to the same scenario, trajectories generated by each domain-specific expert are drawn in different colors.

The first example (Fig. 6(a)) depicts a scenario where the ego vehicle passes through an intersection, with expert behaviors categorized into either left-turning or proceeding straight. The trajectory processed through routing network fusion prioritizes the straight-ahead solution from the orange expert in its planned trajectory.

Fig. 6(b), as the ego vehicle approaches a T-intersection, the yellow expert opts for a right turn while the remaining experts

Table I. Performance on the Navtest Benchmark with Original Metrics.

Type	Method	Input	NC $\uparrow$	DAC $\uparrow$	EP $\uparrow$	TTC $\uparrow$	C $\uparrow$	PDMS $\uparrow$
Diffusion-based	DiffusionDrive [17]	C & L	<u>98.2</u>	<b>96.2</b>	<b>82.2</b>	<u>94.7</u>	<b>100</b>	<b>88.1</b>
IL-based	ADMLP [24]	-	93.0	77.3	62.8	83.6	<b>100</b>	65.6
	VADv2 [14]	C & L	97.2	89.1	76.0	91.6	<b>100</b>	80.9
	UniAD [4]	C	97.8	91.9	78.8	92.9	<b>100</b>	83.4
	LTF [13]	C	97.4	92.8	79.0	92.4	<b>100</b>	83.8
	Transfuser [13]	C & L	97.7	92.8	79.2	92.8	<b>100</b>	84.0
	PARA-Drive [25]	C	97.9	92.4	79.3	93.0	99.8	84.0
	DRAMA [12]	C & L	98.0	93.1	80.1	<b>94.8</b>	<b>100</b>	85.5
	GoalFlow* [26]	C & L	<b>98.3</b>	93.8	79.8	94.3	<b>100</b>	85.7
Hydra-MDP* [15]	C & L	<b>98.3</b>	<u>96.0</u>	78.7	94.6	<b>100</b>	86.5	
Hydra-MDP++* [16]	C	97.6	<u>96.0</u>	80.4	93.1	<b>100</b>	86.6	
IL-based	<b>Ours</b>	C & L	<b>98.3</b>	95.1	<u>81.4</u>	94.2	<b>100</b>	<u>86.9</u>

In the ‘‘Input’’ column, ‘‘C’’ denotes the use of camera images as sensor input, and ‘‘C&L’’ represents the use of both camera images and lidar as sensor inputs. All models require ego-vehicle state as input. \* For fair comparison, we used the official scores of versions with the same backbone network and hidden layer dimensions. The best and second-best scores are highlighted in **bold** and underlined, respectively.

Table II. Performance on the Navtest Benchmark with Extended Metrics.

Method	NC $\uparrow$	DAC $\uparrow$	EP $\uparrow$	TTC $\uparrow$	TL $\uparrow$	DDC $\uparrow$	LK $\uparrow$	HC $\uparrow$	EC $\uparrow$	EPDMS $\uparrow$
Ego Statue MLP	93.1	77.9	86.0	91.5	99.6	92.7	89.4	98.3	85.4	64.0
Transfuser [13]	96.9	89.9	<u>87.1</u>	95.4	<u>99.7</u>	97.8	92.7	98.3	<u>87.2</u>	76.7
Hydra-MDP++ [16]	<u>97.2</u>	<b>97.5</b>	83.1	<u>96.5</u>	<b>99.8</b>	<b>99.4</b>	<u>94.4</u>	<u>98.2</u>	70.9	<u>81.4</u>
Ours	<b>98.3</b>	<u>95.1</u>	<b>87.6</b>	<b>97.4</b>	<b>99.8</b>	<u>98.6</u>	<b>96.5</b>	<b>98.3</b>	<b>89.1</b>	<b>83.1</b>

choose to continue straight ahead. The trajectory processed through routing network fusion incorporates elements from the yellow expert’s result.

Fig.6(c) and Fig.6(d) further illustrate two different scenarios: navigating a roundabout and lane entry selection. In the Fig.6(c) case, the experts’ planning proposals show significant divergence: except for the red expert, the others incorrectly handle the scenario by simply choosing to drive toward the right-front direction, whereas the red expert successfully captures the roundabout road features and correctly navigates around it. Notably, the final planning trajectory using the routing network for expert fusion primarily references the red expert’s solution, demonstrating the effectiveness and rationality of the intrinsic routing mechanism. The visualization of the lane entry selection scenario (Fig.6(d)) also clearly reveals different experts’ preferences for different entrance lane positions in the current context.

In addition, we also visually compared the trajectory planning performance of ARTEMIS and Transfuser across different scenarios, with results shown in Fig.7.

#### E. Ablation Study

**Component-Wise Impact Analysis.** To assess the contribution of each architectural element, we constructed three ablated variants by individually removing the Autoregressive Planning Module with Mixture-of-Experts (AME), the Mixture-of-Experts module (MOE), and the trajectory refinement module (TR). The results, summarized in Tab.III, show that eliminating any single component degrades performance, thereby confirming the necessity of all three. Specifically, omitting the autoregressive module reduces PDMS by 2.9 points, indicating that the autoregressive paradigm is essential for capturing temporal dependencies among waypoints and for accurately accommodating evolving environmental context. Removing the MoE module yields a PDMS drop of 4.0

points, underscoring the advantage of the MoE architecture in dynamically adapting to diverse driving scenarios and behavior patterns. Finally, excising the cascaded-refinement module lowers PDMS by 2.2 points, demonstrating that this stage effectively mitigates sampling instabilities that can arise during autoregressive trajectory generation.

Table III. Ablation Study of Model Components.

AME	MOE	TR	NC	DAC	EP	TTC	C	PDMS
×	✓	✓	97.8	92.3	78.9	93.4	99.9	84.0
✓	×	✓	97.6	91.7	78.2	92.7	99.8	82.9
✓	✓	×	97.7	93.1	79.7	93.3	99.9	84.7
✓	✓	✓	<b>98.3</b>	<b>95.1</b>	<b>81.4</b>	<b>94.2</b>	<b>100</b>	<b>86.9</b>

**The Routing Network.** To verify the effectiveness of the intrinsic routing mechanism, we compare it with an explicit driving-command guidance scheme. The experimental results in Tab.IV show that the intrinsic routing mechanism effectively prevents inappropriate assignment of domain-specific experts when explicit guidance deviates from the actual scene.

We also conducted experiments by fixing the activation to single experts within the model. The results presented in Table.V demonstrate that activating any individual expert in isolation fails to achieve optimal performance, underscoring the necessity of the routing network for facilitating effective expert allocation. Moreover, the results confirms that the model training process did not exhibit disproportionate preference toward utilizing any particular expert, indicating balanced expert utilization throughout the learning process.

Table IV. Comparison of Expert-Assignment Strategies.

Expert Assignment	NC	DAC	EP	TTC	C	PDMS
Driving Cmd.	97.8	92.1	78.5	93.0	99.8	83.5
Intrinsic Routing	<b>98.3</b>	<b>95.1</b>	<b>81.4</b>	<b>94.2</b>	<b>100</b>	<b>86.9</b>

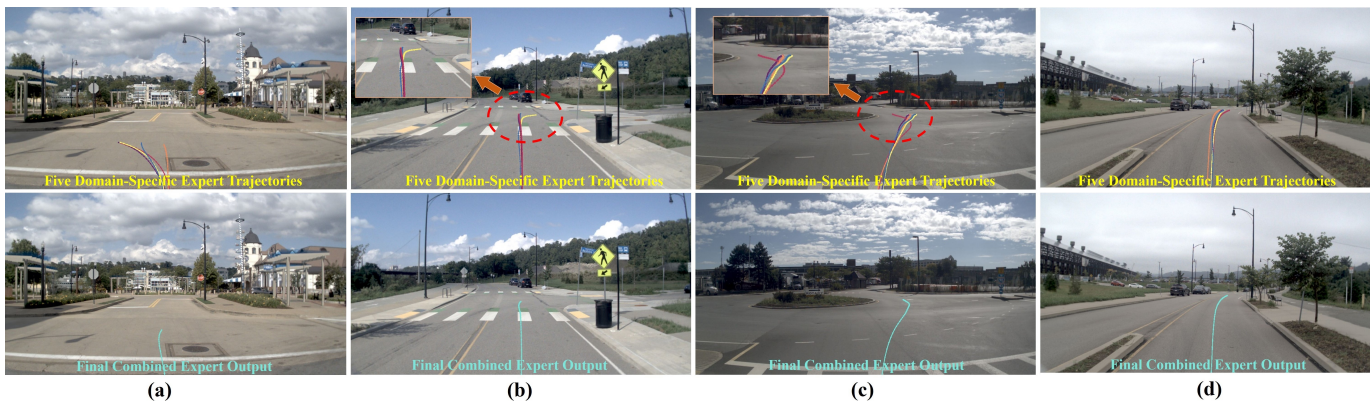


Fig. 6. Qualitative analysis results of ARTEMIS on the navtest split. The qualitative results demonstrate the differentiated processing strategies of five domain-specific experts and the mixture expert for the same scenario. The visualization trajectories are sampled from 2Hz to 10Hz.

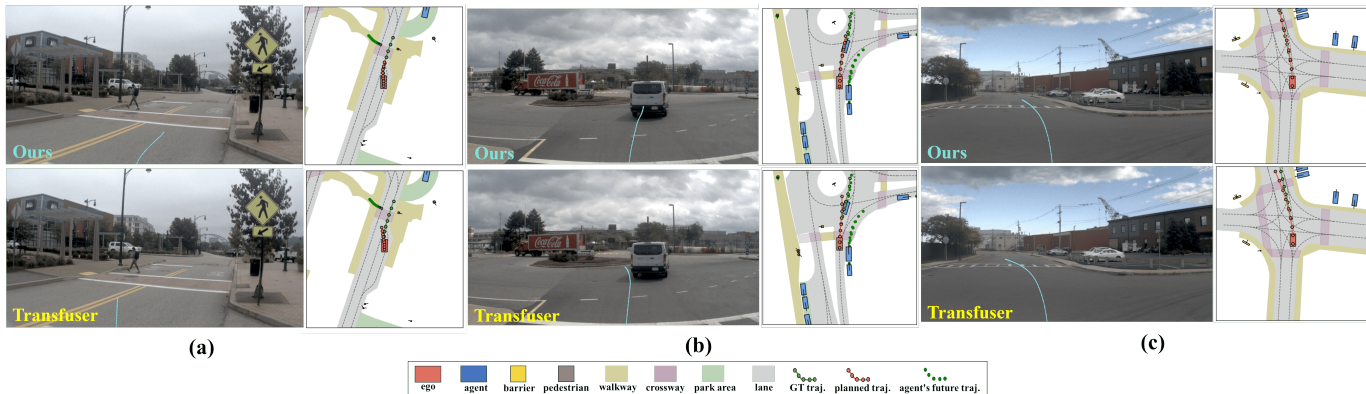


Fig. 7. Qualitative comparison between ARTEMIS and the baseline Transfuser on the navtest split. The results indicate that ARTEMIS, which employs a mixture of experts model, can more accurately respond to erroneous control commands (Figure (a)), while significantly improving capabilities such as drivable area compliance (Figure (b)) and lane keeping (Figure (c)), demonstrating enhanced trajectory planning stability.

Table V. Performance of Single Expert Activation Strategies

Individually Activated Expert	NC	DAC	EP	TTC	C	PDMS
Expert 1	97.5	92.5	79.0	92.2	99.9	83.4
Expert 2	98.0	94.7	81.0	93.7	100	86.2
Expert 3	97.4	91.0	77.4	91.9	99.8	81.9
Expert 4	97.8	93.8	80.0	93.3	99.9	85.1
Expert 5	98.1	94.1	80.2	93.2	100	85.4
Intrinsic Routing	<b>98.3</b>	<b>95.1</b>	<b>81.4</b>	<b>94.2</b>	<b>100</b>	<b>86.9</b>

**Cascaded Refinement Layers.** Tab.VI reports how varying the number of refinement layers influences model performance.

Table VI. Impact of the Number of Refinement Layers.

Refinement Layers	NC	DAC	EP	TTC	C	PDMS
1	98.1	93.9	80.0	<b>94.2</b>	<b>100</b>	85.7
2	<b>98.3</b>	<b>95.1</b>	<b>81.4</b>	<b>94.2</b>	<b>100</b>	<b>86.9</b>
3	97.9	93.7	80.6	93.2	99.9	85.4

**Domain-Specific Experts.** Tab.VII investigates the impact of varying the number of experts. Enlarging the expert pool from three to five gradually increases the model’s performance. However, expanding to ten experts leads to a performance drop of 1.4 points, suggesting that, given limited training data, an excessive number of experts disperses resources and causes functional overlap.

Table VII. Effect of the Number of Domain-Specific Experts.

Domain-Specific Experts	NC	DAC	EP	TTC	C	PDMS
3	98.2	94.1	80.5	94.1	99.9	85.9
5	<b>98.3</b>	<b>95.1</b>	<b>81.4</b>	<b>94.2</b>	<b>100</b>	<b>86.9</b>
10	98.1	93.6	80.0	93.9	<b>100</b>	85.5

**Multi-Seed Variance Analysis.** To investigate the model’s performance variation under different random seeds, we conducted a multi-seed variance analysis. We trained and evaluated the ARTEMIS model using four different random seeds. As shown in Tab.VIII, the model’s performance across various metrics shows minimal fluctuation between runs, with the standard deviation remaining at a very low level.

Table VIII. Performance Stability Across Different Random Seeds. Results Are Reported as Mean  $\pm$  Standard Deviation.

Random Seeds	NC	DAC	EP	TTC	C	PDMS
0	98.29 $\pm$ 0.004	95.06 $\pm$ 0.005	81.40 $\pm$ 0.006	94.24 $\pm$ 0.003	99.99 $\pm$ 0	86.94 $\pm$ 0.005
1	98.27 $\pm$ 0.005	95.05 $\pm$ 0.004	81.41 $\pm$ 0.005	94.23 $\pm$ 0.005	100.00 $\pm$ 0	86.93 $\pm$ 0.004
2	98.31 $\pm$ 0.003	95.07 $\pm$ 0.005	81.37 $\pm$ 0.004	94.20 $\pm$ 0.002	99.98 $\pm$ 0	86.96 $\pm$ 0.004
3	98.25 $\pm$ 0.004	95.06 $\pm$ 0.003	81.35 $\pm$ 0.005	94.23 $\pm$ 0.002	99.99 $\pm$ 0	86.92 $\pm$ 0.003

## IV. CONCLUSION

This paper introduces ARTEMIS, in contrast to conventional static paradigms that synthesize complete trajectories in one-shot inference paradigms, it implements a sequential decision-making process that enables modeling of trajectory evolution.

Table IX. Performance on the NAVSIM v2 Benchmark.

Method	Stage	NC $\uparrow$	DAC $\uparrow$	EP $\uparrow$	TTC $\uparrow$	TL $\uparrow$	DDC $\uparrow$	LK $\uparrow$	EPDMS $\uparrow$
MLP	Stage one	86.7	76.4	84.2	91.4	<b>98.5</b>	89.6	92.1	18.1
	Stage two	70.8	64.9	83.5	65.9	<b>97.5</b>	83.0	47.3	
LTF [13]	Stage one	88.6	81.4	86.9	<b>92.8</b>	97.9	<b>99.6</b>	92.1	24.7
	Stage two	78.9	76.1	85.7	75.9	96.4	<b>91.3</b>	50.4	
Ours	Stage one	<b>89.6</b>	<b>84.3</b>	<b>87.4</b>	90.7	97.9	<b>99.6</b>	<b>93.6</b>	<b>29.0</b>
	Stage two	<b>80.8</b>	<b>77.4</b>	<b>87.0</b>	<b>76.1</b>	96.9	91.0	<b>52.6</b>	

Through its integrated Mixture-of-Experts architecture with dedicated routing networks, ARTEMIS dynamically captures the intrinsic dynamic characteristics of driving behavior and effectively accommodates diverse driving environments. Extensive quantitative evaluations conducted on the NAVSIM benchmark demonstrate that ARTEMIS achieves highly competitive performance. Given its flexibility and adaptability, this framework exhibits substantial deployment potential in complex scenarios and establishes a promising direction for future autonomous driving research.

## V. APPENDIX

**Additional Results on NAVSIM v2 Benchmark.** To conduct additional testing, we exclusively utilize image input. The NAVSIM v2 benchmark introduces a pseudo closed-loop evaluation methodology that approximates real closed-loop performance through a two-stage aggregation mechanism. A more detailed description and calculation of the metrics can be found at [27]. The test results and some metrics are shown in Tab. IX. Compared to the baseline LTF, our method demonstrates an improvement of 17% in EPDMS.

## REFERENCES

- Z. Li, W. Wang, H. Li, E. Xie, C. Sima, T. Lu, Q. Yu, and J. Dai, "Bevformer: learning bird's-eye-view representation from lidar-camera via spatiotemporal transformers," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- G. Xin, D. Chu, L. Lu, Z. Deng, Y. Lu, and X. Wu, "Multi-agent trajectory prediction with difficulty-guided feature enhancement network," *IEEE Robotics and Automation Letters*, 2025.
- J. Cheng, Y. Chen, and Q. Chen, "Pluto: Pushing the limit of imitation learning-based planning for autonomous driving," *arXiv preprint arXiv:2404.14327*, 2024.
- Y. Hu, J. Yang, L. Chen, K. Li, C. Sima, X. Zhu, S. Chai, S. Du, T. Lin, W. Wang, *et al.*, "Planning-oriented autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 17 853–17 862.
- B. Jiang, S. Chen, Q. Xu, B. Liao, J. Chen, H. Zhou, Q. Zhang, W. Liu, C. Huang, and X. Wang, "Vad: Vectorized scene representation for efficient autonomous driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 8340–8350.
- P. Wu, X. Jia, L. Chen, J. Yan, H. Li, and Y. Qiao, "Trajectory-guided control prediction for end-to-end autonomous driving: A simple yet strong baseline," *Advances in Neural Information Processing Systems*, vol. 35, pp. 6119–6132, 2022.
- Y. Liu, J. Zhang, L. Fang, Q. Jiang, and B. Zhou, "Multimodal motion prediction with stacked transformers," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 7577–7586.
- Y. Yuan, X. Weng, Y. Ou, and K. M. Kitani, "Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 9813–9823.
- N. Rhinehart, K. M. Kitani, and P. Vernaza, "R2p2: A reparameterized pushforward policy for diverse, precise generative path forecasting," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 772–788.
- Y. Chen, Y. Wang, and Z. Zhang, "Drivinggpt: Unifying driving world modeling and planning with multi-modal autoregressive transformers," *arXiv preprint arXiv:2412.18607*, 2024.
- X. Hu, W. Yin, M. Jia, J. Deng, X. Guo, Q. Zhang, X. Long, and P. Tan, "Drivingworld: Constructing world model for autonomous driving via video gpt," *arXiv preprint arXiv:2412.19505*, 2024.
- C. Yuan, Z. Zhang, J. Sun, S. Sun, Z. Huang, C. D. W. Lee, D. Li, Y. Han, A. Wong, K. P. Tee, *et al.*, "Drama: An efficient end-to-end motion planner for autonomous driving with mamba," *arXiv preprint arXiv:2408.03601*, 2024.
- K. Chitta, A. Prakash, B. Jaeger, Z. Yu, K. Renz, and A. Geiger, "Transfuser: Imitation with transformer-based sensor fusion for autonomous driving," *IEEE transactions on pattern analysis and machine intelligence*, vol. 45, no. 11, pp. 12 878–12 895, 2022.
- S. Chen, B. Jiang, H. Gao, B. Liao, Q. Xu, Q. Zhang, C. Huang, W. Liu, and X. Wang, "Vadv2: End-to-end vectorized autonomous driving via probabilistic planning," *arXiv preprint arXiv:2402.13243*, 2024.
- Z. Li, K. Li, S. Wang, S. Lan, Z. Yu, Y. Ji, Z. Li, Z. Zhu, J. Kautz, Z. Wu, *et al.*, "Hydra-mdp: End-to-end multimodal planning with multi-target hydra-distillation," *arXiv preprint arXiv:2406.06978*, 2024.
- K. Li, Z. Li, S. Lan, Y. Xie, Z. Zhang, J. Liu, Z. Wu, Z. Yu, and J. M. Alvarez, "Hydra-mdp++: Advancing end-to-end driving via expert-guided hydra-distillation," *arXiv preprint arXiv:2503.12820*, 2025.
- B. Liao, S. Chen, H. Yin, B. Jiang, C. Wang, S. Yan, X. Zhang, X. Li, Y. Zhang, Q. Zhang, *et al.*, "Diffusiondrive: Truncated diffusion model for end-to-end autonomous driving," *arXiv preprint arXiv:2411.15139*, 2024.
- R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton, "Adaptive mixtures of local experts," *Neural computation*, vol. 3, no. 1, pp. 79–87, 1991.
- D. Dai, C. Deng, C. Zhao, R. Xu, H. Gao, D. Chen, J. Li, W. Zeng, X. Yu, Y. Wu, *et al.*, "Deepseekmoe: Towards ultimate expert specialization in mixture-of-experts language models," *arXiv preprint arXiv:2401.06066*, 2024.
- J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, *et al.*, "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- D. Dauner, M. Hallgarten, T. Li, X. Weng, Z. Huang, Z. Yang, H. Li, I. Gilitschenski, B. Ivanovic, M. Pavone, *et al.*, "Navsim: Data-driven non-reactive autonomous vehicle simulation and benchmarking," *Advances in Neural Information Processing Systems*, vol. 37, pp. 28 706–28 719, 2024.
- Z. Li, Z. Yu, S. Lan, J. Li, J. Kautz, T. Lu, and J. M. Alvarez, "Is ego status all you need for open-loop end-to-end autonomous driving?" in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 14 864–14 873.
- D. Chu, R. Wang, J. Wang, Q. Hua, L. Lu, and C. Wu, "End-to-end autonomous driving: Advancements and challenges," *China Journal of Highway and Transport*, vol. 37, no. 10, pp. 209–232, 2024.
- J.-T. Zhai, Z. Feng, J. Du, Y. Mao, J.-J. Liu, Z. Tan, Y. Zhang, X. Ye, and J. Wang, "Rethinking the open-loop evaluation of end-to-end autonomous driving in nusences," *arXiv preprint arXiv:2305.10430*, 2023.
- X. Weng, B. Ivanovic, Y. Wang, Y. Wang, and M. Pavone, "Paradrive: Parallelized architecture for real-time autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 15 449–15 458.
- Z. Xing, X. Zhang, Y. Hu, B. Jiang, T. He, Q. Zhang, X. Long, and W. Yin, "Goalflow: Goal-driven flow matching for multimodal trajectories generation in end-to-end autonomous driving," *arXiv preprint arXiv:2503.05689*, 2025.
- AGC2025. (2025) Navsim v2 end-to-end driving challenge 2025. AGC2025. [Online]. Available: (<https://huggingface.co/spaces/AGC2025/e2e-driving-2025>)