CAFE-AD: Cross-Scenario Adaptive Feature Enhancement for Trajectory Planning in Autonomous Driving

Junrui Zhang[1,†], Chenjie Wang[2,†], Jie Peng[3], Haoyu Li[4], Jianmin Ji[2,4,*], Yu Zhang[2,4] and Yanyong Zhang[2,3]

Abstract— Imitation learning based planning tasks on the nuPlan dataset have gained great interest due to their potential to generate human-like driving behaviors. However, open-loop training on the nuPlan dataset tends to cause causal confusion during closed-loop testing, and the dataset also presents a long-tail distribution of scenarios. These issues introduce challenges for imitation learning. To tackle these problems, we introduce CAFE-AD, a <u>C</u>ross-Scenario <u>A</u>daptive <u>F</u>eature <u>E</u>nhancement for Trajectory Planning in <u>A</u>utonomous <u>D</u>riving method, designed to enhance feature representation across various scenario types. We develop an adaptive feature pruning module that ranks feature importance to capture the most relevant information while reducing the interference of noisy information during training. Moreover, we propose a cross-scenario feature interpolation module that enhances scenario information to introduce diversity, enabling the network to alleviate overfitting in dominant scenarios. We evaluate our method CAFE-AD, on the challenging public nuPlan Test14-Hard closed-loop simulation benchmark. The results demonstrate that CAFE-AD outperforms state-of-the-art methods including rule-based and hybrid planners, and exhibits the potential in mitigating the impact of long-tail distribution within the dataset. Additionally, we further validate its effectiveness in real-world environments. The code and models will be made available at https://github.com/AlniyatRui/CAFE-AD.

I. INTRODUCTION

With the rapid development of artificial intelligence and autonomous driving, imitation-based motion planners are increasingly being trained to generate driving trajectories that closely mimic human behavior using observations of dynamic driving environments [1]–[10]. Many studies focus on the research of planning networks using real-world datasets [11], [12]. The nuPlan dataset [12] is used commonly due to its inclusion of large amounts of real driving data, long temporal trajectories, and support for both open-loop and closed-loop simulation. However, the quality of trajectories generated by directly imitating offline collected trajectories in the nuPlan dataset is often limited by the following two aspects: ① Causal Confusion and ② Long-tail Distribution.

[1] Institute of Advanced Technology, University of Science and Technology of China (USTC), Hefei 230026, China. zhangjunrui@mail.ustc.edu.cn

[2] Institute of Artificial Intelligence, Hefei Comprehensive National Science Center, 230088, China. wangchenjie@iai.ustc.edu.cn

[3] School of Artificial Intelligence and Data Science, USTC, Hefei 230026, China. pengjieb@mail.ustc.edu.cn, yanyongz@ustc.edu.cn

[4] School of Computer Science and Technology, USTC, Hefei 230026, China. lihaoyu31801@mail.ustc.edu.cn, yuzhang @ustc.edu.cn

† Equal contribution to this paper.

∗ Corresponding author. jianmin@ustc.edu.cn
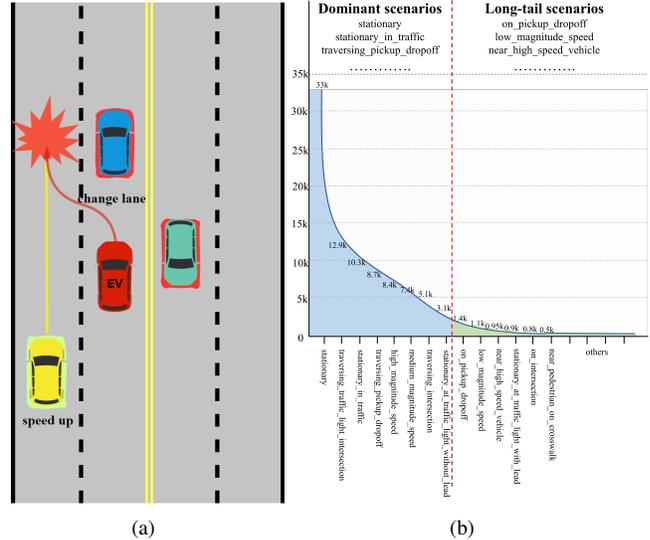
Fig. 1. We illustrate the inherent problems in the dataset. (a) Causal Confusion: Despite the green vehicle being closest to the ego vehicle, it has no causal relationship with the ego's planning. However, the ego vehicle may overly focus on it and neglect the yellow vehicle behind, leading to a potential collision. (b) Long-tail Distribution: Due to the limitations in data collection caused by the complexity of autonomous driving, the distribution across different scenarios is highly imbalanced, which easily leads to the model over-fitting to the dominant scenario, e.g. *stationary*.

a) Causal Confusion: The open-loop training method in nuPlan, which relies on pre-collected trajectories for supervision, tends to cause the model to overly depend on the results of expert behavior. As a result, the model struggles to distinguish causal relationships in the environment, inadvertently focusing on unexpected noise rather than the critical information in the current scenario, as illustrated in Fig. 1(a).

PLUTO [7] uses hand-crafted modules to remove non-reactive agents from future ground truths. It relies on predefined rules, making it less adaptable to dynamic and complex driving environments. Inspired by previous works [13], [14], which employ dynamic token sparsification to allow the network to focus effectively on most informative regions. We aim to further address the issue of causal confusion by grasping the most critical information for the planning process. Specifically, we propose an Adaptive Pruning Transformer Encoder that adaptively identifies the importance of different tokens (e.g., agents, obstacles) using attention probabilities during the encoding phase. It removes tokens and attention connections deemed unimportant by the model, effectively filtering out irrelevant information.

b) Long-tail Distribution: There is a significant disparity in the scale of different scenario types within the

dataset, as illustrated in Fig. 1(b). Dominant scenarios tend to cause the model to overfit their data distribution, leading to poor performance in long-tail scenarios. Previous research [15], [16] improves generalization in model training through mix-up in the feature space. Inspired by these methods, we propose the cross-scenario feature interpolation to address the over-fitting to dominant scenarios caused by the long-tail distribution in the dataset. By introducing a scenario classifier to decompose features into scenario-relevant and scenario-generic components, we perform feature enhancement through the interpolation of scenario-relevant features across different scenario types. Thereby, this module improves the model's generalization ability and robustness across diverse scenarios.

Our proposed CAFE-AD (Fig. 2), Cross-Scenario Adaptive Feature Enhancement for Trajectory Planning in Autonomous Driving, effectively mitigates the interference of the aforementioned challenge, enhancing the planner's ability to capture critical scene information and improving the representation capability across various scenarios. Leveraging our approach, CAFE-AD demonstrates competitive performance compared to current state-of-the-art methods in both closed-loop interactive and non-interactive simulations on the Test14-Hard Benchmark. Additionally, we deploy our designed planner in real-world scenarios on self-driving platforms to generate reference trajectories and further validate its effectiveness.

The primary contributions of our research are outlined as follows:

- We propose CAFE-AD, an effective feature enhancement method designed to improve the performance of imitation learning based planners, ensuring safer and more efficient autonomous driving.
- The design of the Adaptive Pruning Transformer Encoder specifically enhances the model's ability to focus on critical scene elements by selectively pruning irrelevant information. It improves the model's understanding of the causal relationships in the current driving scenario, leading to better performance.
- The cross-scenario feature interpolation is proposed to enhance the information in dominant scenarios with scenario-relevant features from different types. It improves the diversity of scenarios during training, preventing the over-fitting caused by the long-tail distribution.
- Our method demonstrates competitive performance compared to state-of-the-art methods on the Test14-Hard Benchmark. Furthermore, we validated the effectiveness of our approach in real-world scenarios, underscoring its practical value.

II. RELATED WORK

A. Imitation Learning

Imitation learning based motion planners are primarily divided into two categories: end-to-end (E2E) approaches [3], [17]–[21], which directly generate driving trajectories or control signals from raw sensor inputs, and mid-to-mid approaches [9], [22], [23] which utilize processed perception outputs as inputs.

Some E2E approaches [24]–[27] have integrated perception, prediction, and planning into a unified architecture, where each module produces its own output while sharing extracted features with downstream tasks to form an end-to-end output. With the advancement of large language models (LLMs), the recent work [28]–[30] has started to explore the integration of LLMs to leverage their scene understanding and logical reasoning capabilities. However, E2E methods need large amounts of quality data and are hard to debug due to their lack of interpretability, making real-world applications challenging despite validation in simulators like CARLA [31].

This paper primarily focuses on mid-to-mid approaches, which are more modular and interpretable, making them widely applicable to real-world data [1], [4], [32]. The release of the nuPlan dataset [12] has provided recorded real-world training data and a simulator for both open-loop and closed-loop testing. The PDM method [33], which relies solely on the centerline rather than more complex scene representations [4], [6], [34]–[36] achieved first place in the nuPlan challenge 2023. In the same competition, several methods [9], [22], [23] achieved high scores by combining motion planning networks with post-processing optimization techniques. Additionally, other intriguing works include integrating occupancy prediction to guide planning [37], leveraging a Multi-modal LLM to integrate human-like reasoning and common sense [38], the tree-structured policy planning with differentiable joint training for motion prediction and cost evaluation [8], endeavoring to mask the ego vehicle's information [10].

However, pure imitation learning typically suffers from limitations such as causal confusion and long-tail distribution, which result in poor performance of models. To address this, our approach tackles causal confusion by selecting more informative features from the current driving scenario and mitigates the long-tail distribution problem by introducing diverse representations from other scenario types.

B. Enhancement Strategies

Several studies have aimed to overcome the limitations of pure imitation learning by employing augmentation techniques that enhance the learning process, thereby improving the models' applicability and effectiveness in diverse and complex scenarios. The most prevalent enhancement strategies commence at the input data level, employing techniques such as data perturbation [1], [5] and reasonable trajectory augmentation [23] to introduce variability into the learning process. PLUTO [7], PlanTF [5] address causal confusion during training by adding a Dropout layer to input features. Additionally, incorporating supplementary loss functions, for instance, those that penalize collisions and deviations from the designated path [1], [2], [4] is also a critical strategy for enhancing the model's overall efficacy.
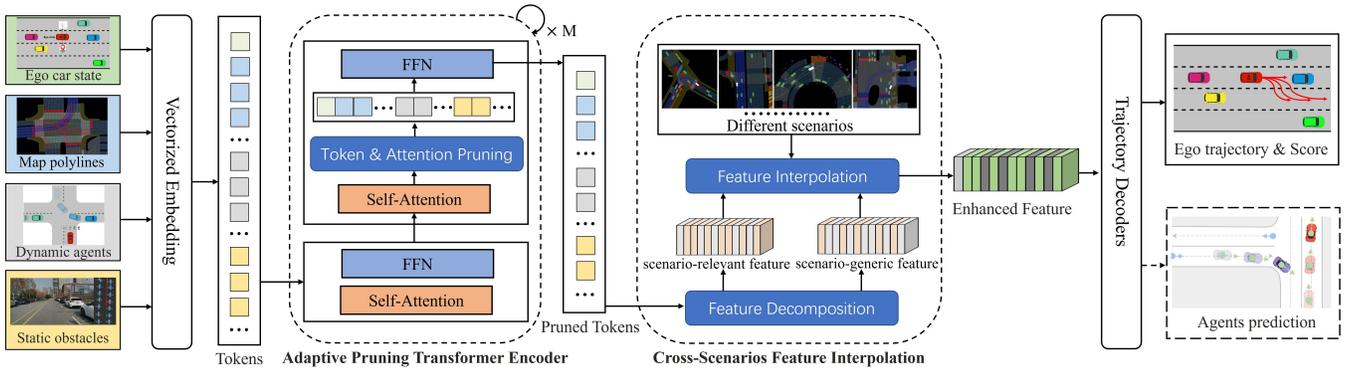
Fig. 2. The overview of our proposed CAFE-AD method.

Compared to these methods that typically perform enhancement in the input space or add auxiliary loss function, our research introduces a novel attempt to achieve more versatile performance through cross-scenario adaptive feature enhancement in the encoded space.

III. METHODOLOGY

In section III-A, we first introduce the model's scene representation input. We then delve into the specifics of the feature enhancement framework, covering the adaptive pruning module in Section III-B and the cross-scenario feature interpolation module in Section III-C. The comprehensive workflow of our research is illustrated in Fig. 2.

A. Scene Representation

In this study, we focus on autonomous driving trajectory planning in urban traffic environments, which involves key elements such as the ego vehicle (EV), surrounding dynamic agents (A), high-definition maps (M), static obstacles (O), and various traffic semantic information. We use PLUTO as the baseline for our method, initially encoding information from different traffic elements into a common feature dimension D, where the encoded vectors are denoted as X_{EV}, X_A, X_M, X_O. To achieve information sharing, we concatenate different vectors into a tensor. To compensate for the loss of global positional information during vectorization, each token is introduced with a Fourier positional embedding X_{PE}, and traffic semantic information is also encoded as learnable embedding X_{LE} and added to the tensor, forming the initial scene understanding embedding X_0 as

$$X_0 = \text{concat}(X_{EV}, X_A, X_M, X_O) + X_{PE} + X_{LE} \quad (1)$$

where the X_0 as input then processed by a multi-layer Adaptive Pruning Transformer Encoder denoted as APT-Encoder, where tokens corresponding to different traffic elements are reordered based on their adaptive importance ranking.

B. Adaptive Pruning Transformer Encoder

Causal confusion arises because the model struggles to retrieve crucial features from the complex driving environment. To address this, we developed an adaptive pruning transformer encoder, applying token- and attention-level pruning to help the model focus on features that are more

important for trajectory planning. After being fed into the APT-Encoder, X_0 is initialized using self-attention [39], as follows. The processed results are then divided into features corresponding to each traffic element in Eq. (1).

$$\text{Self-Attention}(X_0) = [F_{EV}, F_A, F_M, F_O] \quad (2)$$

where, in the self-attention module, the input is passed through three linear transformations to generate the query (Q), key (K), and value (V) matrices. The self-attention matrix A is calculated as follows.

$$A = \text{Softmax}(QK^T/\sqrt{D})V \quad (3)$$

Here we denote the token length of X_0 as L, the query matrix Q and the key matrix K are constructed by concatenating the token representations, where $Q = [q_1, q_2, q_3...q_L]$ and $K = [k_1, k_2, k_3...k_L]$. We use the attention map of the ego-vehicle state token computed using the following equation, to represent the contribution of other tokens to the ego-vehicle's planning performance.

$$a_{EV} = \text{Softmax}(q_{EV}K^T/\sqrt{D}) \quad (4)$$

We utilize $\tilde{a}_{EV}$ after excluding the first element of a_{EV}, as the criterion for selecting tokens. In the multi-head self-attention mechanism, the average attention probabilities $\bar{a}_{EV} = \frac{1}{H}\sum_{i=1}^{H}\tilde{a}_{EV}^{(h)}$ across all heads are used as the basis for ranking and selecting tokens. Therefore, for token pruning, based on the values of $\bar{a}_{EV}$, we select a predefined number of the most important tokens from F_A, F_M and F_O according to a specified ratio π_p. These selected tokens are then concatenated to obtain pruned feature F (Eq. (5)). In attention-level pruning, we select specific attention connections. Specifically, each token computes attention only with the selected tokens.

$$F = \text{concat}(F_{EV}, \text{select}(F_A), \text{select}(F_M), \text{select}(F_O)) \quad (5)$$

After processing by the APT-Encoder, the pruned feature with more relevant information is obtained and fed into the Cross-Scenario Feature Interpolation module denoted as CSFI for diversity enhancement.

C. Cross-Scenario Feature Interpolation

While the pruning mechanism helps the model better capture crucial intra-scenario information, it does not address the

optimization challenge caused by the long-tail distribution of the overall dataset. To mitigate the model over-fitting to the dominant scenario, we developed a cross-scenario feature interpolation method. We initially use a scenario classifier to predict the scenario category based on the encoded features F. The cross-entropy loss $\mathcal{L}_{\text{CE}}$ is calculated between the predicted scenario category and the ground truth. The gradient of the loss with respect to the encoded features F is denoted as $\nabla_F \mathcal{L}_{\text{CE}}$. We compute the contribution of each feature dimension i to the scenario classifier's predicted logit C_i as:

$$C_i = \text{Mean}(F_i \cdot \nabla_{F_i} \mathcal{L}_{CE}, dim = 0) \quad (6)$$

Then we determine the threshold τ based on the ratio π_o by calculating the quantile of the contributions.

$$\tau = \text{Quantile}([C_0, C_1, \ldots, C_D], \pi_o) \quad (7)$$

where the threshold τ is used to decompose the features along the embedding dimension into two components, as shown below:

$$F = \mathbb{I}(C_i > \tau) \odot F + \mathbb{I}(C_i \leq \tau) \odot F \quad (8)$$

where the former represents features with higher relevance to specific scenarios, denoted as scenario-relevant features F_r, while the latter represents scenario-generic features F_g that are shared across different scenarios. Here, $\odot$ denotes the element-wise multiplication.

We enhance the scenario-relevant feature F_r for introducing further diversity through interpolation with scenario-relevant features F_r' from different scenarios. Here π_r represents the interpolation ratio, which is sampled from a uniform distribution $\mathcal{U}(0,1)$. Note that only the dominant scenario is augmented. In our experiments, dominant scenarios are defined as those whose quantities exceed the average number per scenario. The process is shown below:

$$F_{interpolation} = F_g + (1 - \pi_r) \times F_r + \pi_r \times F_r' \quad (9)$$

Finally, the enhanced features are passed through a Transformer decoder to generate the ego vehicle's future trajectory and predict the behavior of dynamic agents. By incorporating features from different scenarios, we introduce diversity to encourage the model to balance fitting across scenarios.

D. Training Process

We denote ϕ as the feature decoder. The loss function takes the feature E and the decoder's output $\phi(E)$ as inputs, following the design of PLUTO [7], which incorporates contrastive learning loss, auxiliary loss, ego-trajectory imitation loss, and surrounding agents' prediction loss.

$$\mathcal{L}(E, \phi(E)) = \mathcal{L}_{cll} + \mathcal{L}_{aux} + \mathcal{L}_{ego} + \mathcal{L}_{agents} \quad (10)$$

We train the model in two phases. During the warm-up phase, the planning model is trained on the original encoded feature E following $\mathcal{L}(E, \phi(E))$. The scenario category classifier θ is trained using E to predict the scenario category labels i with $\mathcal{L}_{CE}(\theta(E), i)$.

After the warm-up phase is completed, we use both the original and augmented features with the following objective function:

$$\mathcal{L}_{aug} = \frac{1}{2} \left(\mathcal{L}(E, \phi(E)) + \mathcal{L}(E', \phi(E')) \right) \quad (11)$$

where the E' represents the augmented feature from the original feature E, notice that due to the pruning process, the agents are reordered, thus the prediction ground truth for $\mathcal{L}_{agents}$ in Equation 10 is modified correspondingly. Meanwhile, the scenario category classifier uses the pruned token to compute the loss with $\mathcal{L}_{CE}$ and updates accordingly, where the token does not use interpolation to change the distribution but simply drops a few insignificant features.

IV. EXPERIMENTAL SETUP

A. Implementation Details

In this work, we propose an augmentation method to enhance the driving performance of imitation learning planners. Our prune ratio π_p is set to 0.9 and is applied every two layers of the Transformer Encoder. We implement a cyclic threshold adjustment scheme for the scenario-relevant threshold π_o that enables the model to adapt incrementally to increasing enhancement strength. The π_o is initiated with a threshold of 0.9, which is decremented by 0.1 every 100 steps, reaching a minimum of 0.5 before resetting to 0.9 to repeat the process. In accordance with the nuPlan challenge 2023 design, we use 2 seconds of historical data to predict 8 seconds of future trajectory. The model is trained on 4 NVIDIA A30 GPUs with a batch size of 128 for 25 epochs. We use the AdamW optimizer with a weight decay set to $1e^{-4}$. The learning rate is linearly increased to $1e^{-3}$ over the initial three warm-up epochs, followed by a cosine decay schedule for the rest of the training.

B. Evaluation Setting

For all experiments, we use a consistent training split of 100k frames sampled from all scenario types for training. For evaluation, we use the nuPlan Test14-Hard Benchmark which is challenging for imitation learning based planners. As detailed in PlanTF [5], the Test14-Hard Benchmark was created by running 100 scenarios for each of the 14 scenario types specified in the nuPlan challenge 2023 using a powerful rule-based planner (PDM-Closed) and then selecting the 20 least-performing scenarios for each type. The nuPlan dataset consists of three simulation challenges: open-loop simulation, non-reactive closed-loop simulation, and reactive closed-loop simulation. Previous work [40] has demonstrated a significant gap between the open-loop evaluation metrics and real driving performance, so we only conduct closed-loop simulations. The evaluation metrics for closed-loop simulations include: (1) Traffic Rule Compliance: The evaluation of the EV's adherence to speed limits, maintenance of the correct driving direction, and the ability to stay within the drivable area boundaries. (2) Safety: The assessment of vehicle safety by checking the avoidance of ego-at-fault collisions and ensuring that the Time to Collision (TTC)

Type	Methods	R-Score	Collisions	TTC	Drivable	Comfort	Progress	Speed	NR-Score
Expert	Log-Replay	68.80	77.02	69.85	95.96	99.26	98.48	94.12	85.96
Rule-based	IDM	62.26	84.38	72.43	84.19	87.87	69.6	96.52	56.16
	PDM-Closed	75.19	95.22	84.19	95.59	83.46	75.48	**99.54**	65.08
Pure-learning	UrbanDriver	36.39	63.05	56.62	72.43	**99.63**	**79.99**	79.19	35.71
	DTPP	42.67	88.6	84.93	94.12	94.12	38.26	99.28	41.60
	PlanTF	58.76	89.52	83.82	91.91	83.82	65.38	97.79	60.53
	PLUTO (w/o post.)	58.74	87.68	81.62	94.12	86.03	64.87	97.71	63.24
	CAFE-AD (w/o post.)	65.25	88.6	82.35	94.85	88.6	71.64	97.43	68.84
Hybrid	GameFormer	69.09	92.10	82.72	91.91	91.91	68.73	98.00	66.74
	PLUTO	75.35	95.59	87.87	97.06	78.68	75.07	97.8	73.87
	CAFE-AD (Ours)	**78.16**	**95.96**	**88.97**	**97.43**	83.82	77.51	97.93	**76.04**

exceeds a specified threshold. (3) The evaluation of comfort through acceleration levels and the assessment of progress by comparing route coverage to that of the expert driver.

C. SOTA Methods

We compare current state-of-the-art methods with our model. (1) **IDM** [41]: A classic rule-based method that adjusts vehicle behavior. (2) **PDM-Closed** [33]: An advanced rule-based algorithm that won the nuPlan challenge 2023. It generates multiple candidate trajectories through various hyper-parameters of IDM and ultimately selects the optimal trajectory. (3) **UrbanDriver** [4]: A planner that uses vectorized inputs with polyline encoders based on PointNet and Transformers. Its open-loop variant is employed for evaluation, incorporating history perturbation during training. (4) **GameFormer** [9]: A k-level game-based interactive trajectory prediction and motion planning architecture. (5) **DTPP** [8]: An integrating joint training method for ego-conditioned motion prediction and cost evaluation within a tree policy planner. (6) **PlanTF** [5]: A concise and efficient learning-based baseline based on the Transformer architecture. (7) **PLUTO** [7]: A powerful method based on a longitudinal-lateral aware model architecture, enhanced with auxiliary loss techniques and contrastive learning to further improve performance.

All methods are trained on our 100k-split random datasets, following the configurations provided in their open-source versions, except for DTPP due to the failure in data processing. For DTPP, we ensure that the scenarios in the training set are contained in our split, and we introduce additional data for validation. Notice that both PLUTO and CAFE-AD initially encountered simulation errors in approximately 2% of the Test14-Hard Benchmark. These errors were due to failures in the post-processing in some frames, so we used pure learning results to prevent these errors in those frames. Our modifications will be released along with the code.

V. EXPERIMENTS

A. Quantitative Results

The comparative analysis with current state-of-the-art methods on the Test14-hard benchmark is detailed in Table I. Our proposed CAFE-AD significantly outperforms all other methods across both evaluation tasks. Compared to

the current leading method, PLUTO, our approach achieves a 2.81 score improvement in CLS-R and a 2.17 score improvement in CLS-NR. Moreover, when considering the pure learning variant without post-processing (w/o post.), we further demonstrate a more significant advantage over PLUTO's variant with a lead of 6.51 and 5.6 in score on the two tasks respectively. Notably, we surpass PLUTO across nearly all metrics, regardless of whether post-processing is applied. In particular, we show a great improvement in the $Progress$ metric. This is primarily because we mitigate over-fitting to dominant scenarios like $stationary$, allowing the pure learning approach to exhibit more proactive driving behavior without relying on rule-based guidance.

B. Case Analysis

Fig. 3 presents a qualitative analysis of the driving behavior between our method and the baseline PLUTO in the typical scenario from the Test14-hard benchmark. It also shows the changes in normalized attention weights toward surrounding agents at different driving moments. As shown in the figure, our method consistently identifies the $vehicle\_01$ ahead as the most critical agent (with a normalized attention weight of 1). In contrast, PLUTO alternates its focus between $vehicle\_01$ and $vehicle\_02$. This causes PLUTO to fail in making a timely and safe interaction with vehicle $vehicle\_01$, resulting in a collision at Frame 45. In comparison, our method uses the APT-Encoder to focus on the most important information in the current driving scene, reducing the impact of causal confusion. It begins deceleration at Frame 20, successfully avoiding a collision with $vehicle\_01$ which has the most significant impact on the current driving behavior, ensuring safe driving.

Fig. 4 illustrates how our method alleviates the issue of over-fitting to dominant scenarios caused by the long-tail distribution. In nuPlan dataset, approximately 45% of the scenarios are static, e.g. $stationary$, $stationary\_in\_traffic$. As shown in the figure, this leads the PLUTO method to remain stationary in some scenarios where expert trajectories suggest moving forward, due to over-fitting to static scenarios. In contrast, our method continues to exhibit proactive driving behavior, demonstrating the effectiveness of our cross-scenario feature interpolation.
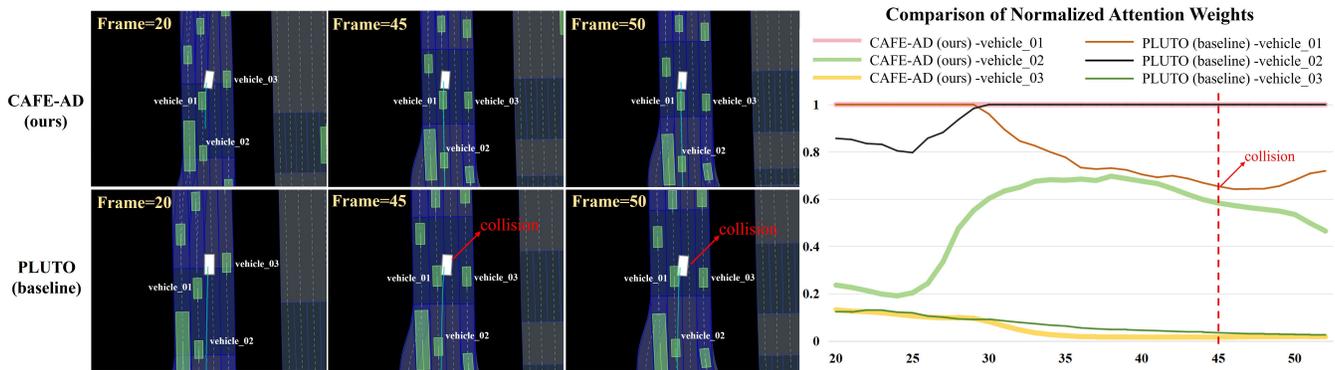
Fig. 3. Qualitative evaluation between CAFE-AD (ours) and PLUTO (baseline) and the changes in normalized attention weights.
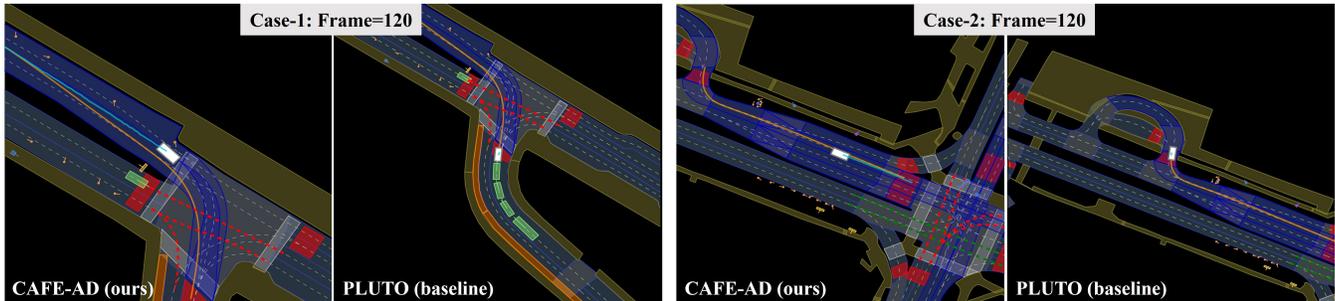

Fig. 4. Qualitative evaluation of handling the long-tail distribution problem.

TABLE II

ABLATION STUDY OF MODULE DESIGN.

APT	CSFI	R-Score	Collisions	Comfort	Progress	NR-Score
		58.74	87.68	86.03	64.87	63.24
√		64.24	**91.36**	84.56	64.01	65.2
	√	64.95	88.79	86.03	67.0	66.73
√	√	**65.25**	88.6	**88.6**	**71.64**	**68.84**

TABLE III

ABLATION STUDY OF DIFFERENT BASELINE.

Method	R-Score	Collisions	Comfort	Progress	NR-Score
PlanTF	58.76	**89.52**	83.82	65.38	60.53
PlanTF w. CAFE-AD	**61.29**	87.87	**86.03**	**70.29**	**62.4**

C. Ablation Study

(1) **Ablation Study of Module Design**. Our method incorporates two crucial modules: Adaptive Pruning Transformer Encoder (APT-Encoder) and Cross-Scenario Feature Interpolation (CSFI). We conducted ablation studies to validate the effectiveness of our design, as shown in Table II. It demonstrates that using either the APT-Encoder or CSFI module individually results in improvements over the baseline method. The combination of these modules achieves the best overall performance. This further demonstrates that the combined enhancement from these two modules, through selecting important information within the current scenario and introducing diversity across scenarios, significantly improves planning effectiveness.

(2) **Ablation Study of Different Baseline**. Table III demonstrates the enhancement effect of our method on another baseline-PlanTF. The results show improvements across nearly all metrics, demonstrating the generalization and transferability of the proposed method.

D. Real-World Validation

To demonstrate the effectiveness of our proposed method in real-world applications, we trained the model using expert data collected from an actual vehicle and deployed it on an autonomous driving platform. The model was integrated to provide a reference trajectory for a traditional optimization-based planning method. This hybrid method further illustrates the method's effectiveness in handling interactions with surrounding agents and enhancing safe driving. A demonstration of the real-world scenario can be seen in the attached video.

VI. CONCLUSIONS

In this work, we propose CAFE-AD, a novel cross-scenario adaptive feature enhancement method for trajectory planning in autonomous driving, designed to address causal confusion and mitigate the long-tail distribution issues during the imitation learning training process. CAFE-AD introduces an adaptive pruning transformer encoder that filters out planning-irrelevant information, encouraging the model to focus on critical features. Furthermore, CAFE-AD employs a cross-scenario feature interpolation module to prevent the model from over-fitting to dominant scenarios in the long-tail distribution dataset. Owing to these enhancements, our method achieves competitive and generalizable results on the nuPlan Test14-Hard benchmark and improves the overall driving performance. Additionally, both quantitative and qualitative results demonstrate the effectiveness of the two proposed modules in addressing complex and dynamic driving environments.

References

[1] M. Bansal, A. Krizhevsky, and A. Ogale, "Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst," *arXiv preprint arXiv:1812.03079*, 2018.

[2] J. Zhou, R. Wang, X. Liu, Y. Jiang, S. Jiang, J. Tao, J. Miao, and S. Song, "Exploring imitation learning for autonomous driving with feedback synthesizer and differentiable rasterization," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 1450–1457.

[3] M. Vitelli, Y. Chang, Y. Ye, A. Ferreira, M. Wołczyk, B. Osiński, M. Niendorf, H. Grimmett, Q. Huang, A. Jain, *et al.*, "Safetynet: Safe planning for real-world self-driving vehicles using machine-learned policies," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 897–904.

[4] O. Scheel, L. Bergamini, M. Wolczyk, B. Osiński, and P. Ondruska, "Urban driver: Learning to drive from real-world demonstrations using policy gradients," in *Conference on Robot Learning*. PMLR, 2022, pp. 718–728.

[5] J. Cheng, Y. Chen, X. Mei, B. Yang, B. Li, and M. Liu, "Rethinking imitation-based planners for autonomous driving," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 14 123–14 130.

[6] K. Renz, K. Chitta, O.-B. Mercea, A. Koepke, Z. Akata, and A. Geiger, "Plant: Explainable planning transformers via object-level representations," *arXiv preprint arXiv:2210.14222*, 2022.

[7] 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.

[8] Z. Huang, P. Karkus, B. Ivanovic, Y. Chen, M. Pavone, and C. Lv, "Dtpp: Differentiable joint conditional prediction and cost evaluation for tree policy planning in autonomous driving," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 6806–6812.

[9] Z. Huang, H. Liu, and C. Lv, "Gameformer: Game-theoretic modeling and learning of transformer-based interactive prediction and planning for autonomous driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 3903–3913.

[10] K. Guo, W. Jing, J. Chen, and J. Pan, "Ccil: Context-conditioned imitation learning for urban driving," *arXiv preprint arXiv:2305.02649*, 2023.

[11] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631.

[12] H. Caesar, J. Kabzan, K. S. Tan, W. K. Fong, E. Wolff, A. Lang, L. Fletcher, O. Beijbom, and S. Omari, "nuplan: A closed-loop ml-based planning benchmark for autonomous vehicles," *arXiv preprint arXiv:2106.11810*, 2021.

[13] Y. Rao, Z. Liu, W. Zhao, J. Zhou, and J. Lu, "Dynamic spatial sparsification for efficient vision transformers and convolutional neural networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 9, pp. 10 883–10 897, 2023.

[14] Z. Kong, H. Ma, G. Yuan, M. Sun, Y. Xie, P. Dong, X. Meng, X. Shen, H. Tang, M. Qin, *et al.*, "Peeling the onion: Hierarchical reduction of data redundancy for efficient vision transformer training," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 7, 2023, pp. 8360–8368.

[15] Y. Liu, Y. Zou, R. Qiao, F. Liu, M. L. Lee, and W. Hsu, "Cross-domain feature augmentation for domain generalization," *arXiv preprint arXiv:2405.08586*, 2024.

[16] H. Yao, Y. Wang, S. Li, L. Zhang, W. Liang, J. Zou, and C. Finn, "Improving out-of-distribution robustness via selective augmentation," in *International Conference on Machine Learning*. PMLR, 2022, pp. 25 407–25 437.

[17] F. Codevilla, E. Santana, A. M. López, and A. Gaidon, "Exploring the limitations of behavior cloning for autonomous driving," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9329–9338.

[18] D. Chen, B. Zhou, V. Koltun, and P. Krähenbühl, "Learning by cheating," in *Conference on Robot Learning*. PMLR, 2020, pp. 66–75.

[19] K. Chitta, A. Prakash, B. Jaeger, Z. Yu, K. Renz, and A. Geiger, "Transfuser: Imitation with transformer-based sensor fusion for au-

[20] T. Ye, W. Jing, C. Hu, S. Huang, L. Gao, F. Li, J. Wang, K. Guo, W. Xiao, W. Mao, *et al.*, "Fusionad: Multi-modality fusion for prediction and planning tasks of autonomous driving," *arXiv preprint arXiv:2308.01006*, 2023.

[21] X. Jia, P. Wu, L. Chen, J. Xie, C. He, J. Yan, and H. Li, "Think twice before driving: Towards scalable decoders for end-to-end autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 21 983–21 994.

[22] Y. Hu, K. Li, P. Liang, J. Qian, Z. Yang, H. Zhang, W. Shao, Z. Ding, W. Xu, and Q. Liu, "Imitation with spatial-temporal heatmap: 2nd place solution for nuplan challenge," *arXiv preprint arXiv:2306.15700*, 2023.

[23] W. Xi, L. Shi, and G. Cao, "An imitation learning method with data augmentation and post processing for planning in autonomous driving," *URL https://opendrivelab. com/e2ead/AD23Challenge/Track_4_pegasus_weitao. pdf*, 2023.

[24] 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.

[25] 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.

[26] 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.

[27] D. Chen and P. Krähenbühl, "Learning from all vehicles," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 17 222–17 231.

[28] Z. Xu, Y. Zhang, E. Xie, Z. Zhao, Y. Guo, K.-Y. K. Wong, Z. Li, and H. Zhao, "Drivegpt4: Interpretable end-to-end autonomous driving via large language model," *IEEE Robotics and Automation Letters*, 2024.

[29] W. Wang, J. Xie, C. Hu, H. Zou, J. Fan, W. Tong, Y. Wen, S. Wu, H. Deng, Z. Li, *et al.*, "Drivemlm: Aligning multi-modal large language models with behavioral planning states for autonomous driving," *arXiv preprint arXiv:2312.09245*, 2023.

[30] H. Shao, Y. Hu, L. Wang, G. Song, S. L. Waslander, Y. Liu, and H. Li, "Lmdrive: Closed-loop end-to-end driving with large language models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 15 120–15 130.

[31] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning*. PMLR, 2017, pp. 1–16.

[32] H. Shao, L. Wang, R. Chen, H. Li, and Y. Liu, "Safety-enhanced autonomous driving using interpretable sensor fusion transformer," in *Conference on Robot Learning*. PMLR, 2023, pp. 726–737.

[33] D. Dauner, M. Hallgarten, A. Geiger, and K. Chitta, "Parting with misconceptions about learning-based vehicle motion planning," in *Conference on Robot Learning*. PMLR, 2023, pp. 1268–1281.

[34] N. Hanselmann, K. Renz, K. Chitta, A. Bhattacharyya, and A. Geiger, "King: Generating safety-critical driving scenarios for robust imitation via kinematics gradients," in *European Conference on Computer Vision*. Springer, 2022, pp. 335–352.

[35] S. Pini, C. S. Perone, A. Ahuja, A. S. R. Ferreira, M. Niendorf, and S. Zagoruyko, "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)*. IEEE, 2023, pp. 10 069–10 075.

[36] M. Hallgarten, M. Stoll, and A. Zell, "From prediction to planning with goal conditioned lane graph traversals," in *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2023, pp. 951–958.

[37] H. Liu, Z. Huang, and C. Lv, "Occupancy prediction-guided neural planner for autonomous driving," in *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2023, pp. 4859–4865.

[38] Y. Zheng, Z. Xing, Q. Zhang, B. Jin, P. Li, Y. Zheng, Z. Xia, K. Zhan, X. Lang, Y. Chen, *et al.*, "Planagent: A multi-modal large language agent for closed-loop vehicle motion planning," *arXiv preprint arXiv:2406.01587*, 2024.

tonomous driving," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 11, pp. 12 878–12 895, 2022.

[39] A. Vaswani, "Attention is all you need," *Advances in Neural Information Processing Systems*, 2017.

[40] T. Phong, H. Wu, C. Yu, P. Cai, S. Zheng, and D. Hsu, "What truly matters in trajectory prediction for autonomous driving?" *Advances in Neural Information Processing Systems*, vol. 36, 2024.

[41] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical review E*, vol. 62, no. 2, p. 1805, 2000.