

UniUncer: Unified Dynamic–Static Uncertainty for End-to-End Driving

Yu Gao^{1†*}, Jijun Wang^{2*}, Zongzheng Zhang², Anqing Jiang¹, Yiru Wang¹,
Yuwen Heng¹, Shuo Wang¹, HAO SUN¹, Zhangfeng Hu³, Hao Zhao²

Abstract—End-to-end (E2E) driving has become a cornerstone of both industry deployment and academic research, offering a single learnable pipeline that maps multi-sensor inputs to actions while avoiding hand-engineered modules. However, the reliability of such pipelines strongly depends on how well they handle uncertainty: sensors are noisy, semantics can be ambiguous, and interaction with other road users is inherently stochastic. Uncertainty also appears in multiple forms: classification vs. localization, and, crucially, in both static map elements and dynamic agents. Existing E2E approaches model only static-map uncertainty, leaving planning vulnerable to overconfident and unreliable inputs. We present UniUncer, the first lightweight, unified uncertainty framework that jointly estimates and uses uncertainty for both static and dynamic scene elements inside an E2E planner. Concretely: (1) we convert deterministic heads to probabilistic Laplace regressors that output per-vertex location and scale for vectorized static and dynamic entities; (2) we introduce an uncertainty-fusion module that encodes these parameters and injects them into object/map queries to form uncertainty-aware queries; and (3) we design an uncertainty-aware gate that adaptively modulates reliance on historical inputs (ego status or temporal perception queries) based on current uncertainty levels. The design adds minimal overhead and drops throughput by only ~ 0.5 FPS while remaining plug-and-play for common E2E backbones. On nuScenes (open-loop), UniUncer reduces average L2 trajectory error by 7%. On NavsimV2 (pseudo closed-loop), it improves overall EPDMS by 10.8% and notable stage two gains in challenging, interaction-heavy scenes. Ablations confirm that dynamic-agent uncertainty and the uncertainty-aware gate are both necessary. Qualitatively, UniUncer produces well-calibrated uncertainties that encourage human-like decisions when evidence is unreliable, improving robustness without sacrificing efficiency. Our project page is: <https://bradgers.github.io/Uniuncer>.

I. INTRODUCTION

End-to-end (E2E) autonomous driving models [1]–[36] have emerged as a promising paradigm that directly maps sensor inputs to driving decisions. Compared to modular pipelines that separately process perception, prediction and planning, E2E models jointly optimize all stages within a single framework, resulting in simpler architectures and reduced information loss across modules. This holistic approach has made E2E models increasingly attractive for both academic research and large-scale industrial deployment.

Early works primarily investigated the essential components of an E2E system. UniAD [5] integrated perception, prediction, and planning with explicit tracking and mapping

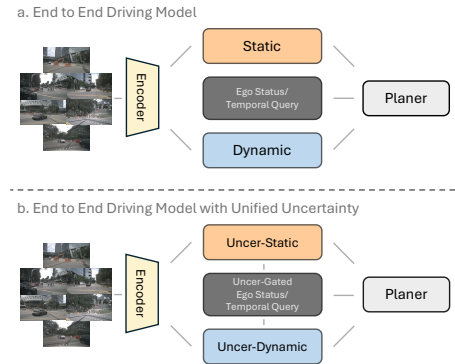


Fig. 1. Comparison between (a) a standard E2E autonomous driving model and (b) our uncertainty-aware E2E model. Our model introduces uncertainty estimation in both static and dynamic branches, and employs an uncertainty-aware gating mechanism to modulate historical information based on current uncertainty levels.

modules, while VAD [6] distilled the design into three core modules: dynamic agents, static map, and planning. Building on these foundations, recent efforts have further expanded the paradigm by incorporating generative modeling [13], [16], [18], world models [29], [37], and large language models (LLMs) [28], [30], [31], thereby enriching scene reasoning, enhancing generalization, and enabling longer-horizon prediction.

Despite these advances, the foundation of E2E systems still lies in the reliable perception of two fundamental elements: static map structures and dynamic agents. However, both elements are inherently uncertain due to noisy sensors, occlusions, ambiguous semantics, and the unpredictable behaviors of surrounding vehicles. While architectural innovations have significantly improved model expressiveness and scalability, explicit modeling of uncertainty, particularly a unified treatment of static and dynamic uncertainties, remains relatively underexplored. This gap is critical, as uncertainty estimation directly influences planning robustness and safety of autonomous systems in complex real-world environments.

Uncertainty has been studied extensively in modular pipelines, particularly in online mapping [36], [38], and UncAD [17] recently extended map-related uncertainty into E2E systems. However, improvements achieved by modeling map uncertainty alone in E2E settings are limited. To fully exploit uncertainty in E2E systems requires, on the one hand, a unified framework that jointly models both dynamic and static uncertainties, and on the other hand, an effective mechanism to incorporate these uncertainties into planning. Nevertheless, uncertainty as a core component remains largely overlooked, and the potential benefits of unified

¹Bosch Corporate Research, China.

²Institute for AI Industry Research (AIR), Tsinghua University, China.

³Rensselaer Polytechnic Institute, USA.

*Equal contribution.

†Corresponding author.

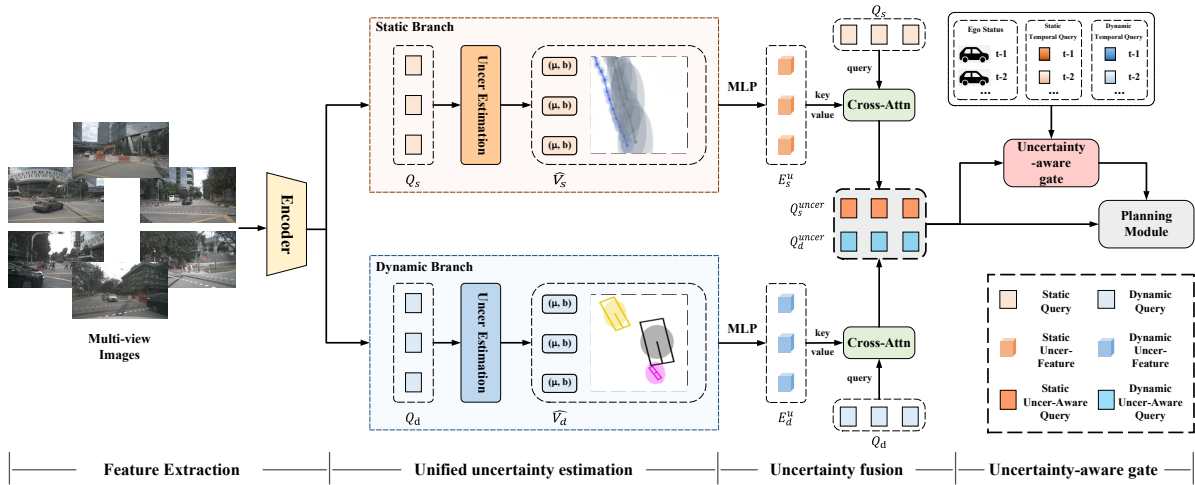


Fig. 2. Overview of the UniUncer framework. UniUncer extends the static and dynamic branches with a unified uncertainty estimation head to predict Laplace parameters. These parameters are encoded and fused with the original queries to form uncertainty-aware queries, which serve as input to the planner and gate historical information (e.g., ego state and temporal queries) via an uncertainty-aware gate.

uncertainty modeling for improving E2E performance are still not well understood. Without explicitly accounting for uncertainties, planning modules may over-rely on unreliable inputs, resulting in degraded robustness and compromised safety in safety-critical driving scenarios.

As illustrated in Fig. 1, a standard E2E model processes all perception features without considering their uncertainty, and integrates all historical information (e.g., past ego status or temporal perception queries) into the planner indiscriminately. In contrast, our uncertainty-aware E2E model in Fig. 1, b explicitly estimates both static and dynamic uncertainties, and uses this information to adaptively gate historical inputs.

In this work, we revisit the role of uncertainty in E2E autonomous driving and propose lightweight modules that explicitly estimate and exploit static and dynamic uncertainties. As shown in Fig. 2, our framework, named UniUncer, enhances planning robustness with minimal architectural modifications and negligible computational overhead. Our main contributions are as follows:

- **Unified uncertainty estimation:** We extend deterministic regression heads in both static and dynamic branches to probabilistic ones, enabling uncertainty-aware feature representations with only minor changes to the existing design.
- **Uncertainty fusion:** We introduce a module that systematically integrates static and dynamic uncertainties into object features, improving downstream planning decisions in challenging scenarios.
- **Uncertainty-aware gate:** We design an adaptive gating mechanism that modulates reliance on historical ego status and temporal perception queries depending on current uncertainty levels, thereby improving the use of historical information for planning decisions.

Extensive experiments on both open-loop (nuScenes) and pseudo closed-loop (NavsimV2) benchmarks demonstrate the effectiveness of our approach. Our method achieves more than 7% reduction in L2 error on nuScenes and a 10%

improvement in EPDMS on NavsimV2, highlighting the benefits of uncertainty-aware modeling for improving the reliability of E2E autonomous driving.

II. RELATED WORKS

A. End-to-End Driving Models

End-to-end (E2E) autonomous driving aims to directly map raw sensor inputs to ego vehicle planning decisions, bypassing the explicit decomposition into modular perception–prediction–planning pipelines. Early approaches such as Transfuser [39] introduced attention-based fusion for multi-modal inputs, while UniAD [5] pioneered a differentiable unified framework that integrates perception, prediction, and planning in a joint optimization process. Subsequent designs have focused on improving efficiency and scene representation: VAD [6] adopted vectorized scene modeling, Para-Drive [14] proposed a parallelized architecture, and SparseDrive [15] introduced a fully sparse design that scales favorably with scene complexity. More recently, generative paradigms [13], [16], [18], world-model-based approaches [29], [37], and large language foundation models [28], [30], [31] have been explored to enhance reasoning, generalization, and long-horizon planning.

Despite these architectural and representational advances, most E2E models treat perception outputs as deterministic and fully reliable, overlooking the inherent uncertainty in both static and dynamic scene elements. While some efforts have begun to incorporate uncertainty, particularly in static map modeling [17], the unified treatment of dynamic and static uncertainties, and their explicit integration into planning, remains largely unaddressed. In contrast, UniUncer introduces a lightweight yet effective framework that jointly estimates and leverages both types of uncertainty to improve planning robustness and adaptability.

B. Evaluation Protocols

Evaluating E2E driving models typically involves open-loop, closed-loop, or hybrid settings. Open-loop evaluation,

TABLE I
OPEN-LOOP E2E PERFORMANCE BENCHMARK ON NUSCENES.

Method	Reference	Input	Img. Backbone	L2 (m) ↓				Col. Rate (%) ↓			
				1s	2s	3s	Avg.	1s	2s	3s	Avg.
ST-P3 [1]	ECCV 2022	Camera	EffNet-b4	1.33	2.11	2.90	2.11	0.23	0.62	1.27	0.71
OccNet [3]	ICCV 2023	Camera	ResNet-50	0.45	0.70	1.04	0.73	0.62	0.58	0.63	0.61
UniAD [5]	CVPR 2023	Camera	ResNet-101	0.45	0.70	1.04	0.73	0.62	0.58	0.63	0.61
VAD [6]	ICCV 2023	Camera	ResNet-50	0.41	0.70	1.05	0.72	0.03	0.19	0.43	0.21
GenAD [13]	ECCV 2024	Camera	ResNet-50	0.36	0.83	1.55	0.91	0.06	0.23	1.00	0.43
DiffusionDrive [16]	CVPR 2025	Camera	ResNet-50	0.27	0.54	0.90	0.57	0.03	0.05	0.16	0.08
BridgeAD [19]	CVPR 2025	Camera	ResNet-50	0.29	0.57	0.92	0.59	0.01	0.05	0.22	0.09
SparseDrive [15]	ICRA 2025	Camera	ResNet-50	0.29	0.58	0.96	0.61	0.01	0.05	0.18	0.08
UncAD [17]	ICRA 2025	Camera	ResNet-50	0.29	0.57	0.95	0.60	0.00	0.04	0.18	0.07
SparseDrive + Ours	-	Camera	ResNet-50	0.27	0.53	0.90	0.57	0.00	0.04	0.17	0.07

widely used on benchmarks such as nuScenes [40] and Navsim [41], assesses trajectory accuracy and rule-based safety metrics. However, it fails to capture long-term feedback loops and multi-agent interactions, limiting its realism. Closed-loop evaluation in simulators like CARLA [42] and Bench2Drive [43] provides more realistic assessment of driving policies but suffers from computational inefficiency and sim-to-real gaps.

To bridge this divide, NavsimV2 [44] introduced a pseudo-simulation framework based on 3D Gaussian Splatting (3DGS) [45], offering a compelling middle ground by combining high-fidelity scene rendering with efficient execution. This enables realistic evaluation without the overhead of full simulation. In our experiments, we adopt both the nuScenes (open-loop) and NavsimV2 (pseudo closed-loop) benchmarks to comprehensively validate the effectiveness of uncertainty-aware modeling under diverse evaluation paradigms.

C. Uncertainty Estimation

Uncertainty estimation plays a critical role in safe and reliable decision-making, particularly in ambiguous or safety-critical environments. In perception tasks, including object detection, segmentation, and depth completion, uncertainty modeling has been shown to improve reliability and interpretability [46]–[49]. For trajectory prediction, incorporating map-related uncertainty has been demonstrated to enhance both accuracy and robustness [36], [38]. More recently, UncAD [17] extended uncertainty modeling to E2E driving by integrating online mapping uncertainty into the planning process.

However, these approaches focus primarily on static map uncertainty and neglect the uncertainty associated with dynamic agents, an equally critical source of risk in interactive driving scenarios. Moreover, they lack a unified mechanism to jointly model and exploit both static and dynamic uncertainties during planning.

III. METHOD

This section begins by introducing UniUncer, a unified uncertainty-aware E2E framework that models both static

and dynamic elements as vectorized Bird’s Eye View (BEV) representations with Laplace-distributed vertices to capture spatial uncertainty. We then describe how a unified uncertainty head estimates these uncertainties (Sec. III-A), how uncertainty-aware queries are generated via uncertainty fusion (Sec. III-B), and how the enhanced queries dynamically gate historical information through an uncertainty-aware gate (Sec. III-C). As shown in Fig. 2, the E2E system captures images from surrounding-view cameras and uses an encoder to produce perception-related queries Q , including static and dynamic queries Q_s and Q_d .

A. Unified Uncertainty Estimation

We model the aleatoric uncertainty of vertex locations with a Laplace distribution. This choice is motivated by its natural connection to the robust ℓ_1 loss commonly used in object detection and vectorized mapping. Formally, minimizing the ℓ_1 loss is equivalent to performing maximum likelihood estimation under a Laplace noise assumption with a fixed scale. By generalizing this to predict both the location μ and the scale b for each vertex, we obtain a principled and learnable measure of spatial uncertainty while maintaining compatibility with existing regression frameworks.

Uncertainty estimation for static objects has proven to be effective in both modular [36], [38] and E2E [17] autonomous driving systems. These approaches explicitly model aleatoric uncertainty, which characterizes the stochastic noise inherent to the measurement process [46]. They extend vectorized map representations by modeling each element as a set of N_s vertices in the BEV space, with each vertex parameterized by a Laplace distribution in the x and y directions:

$$\mathbf{v}_i^s = (\mu_x^{(i)}, b_x^{(i)}, \mu_y^{(i)}, b_y^{(i)}) \in \mathbb{R}^4, \quad i = 1, \dots, N_s,$$

where μ and b denote the location and scale parameters, respectively. The Laplace distribution is a natural choice, as most detection models are trained with regression loss ℓ_1 .

To enable a unified representation of static and dynamic objects, we convert the bounding box labels online during training. The conventional box param-

ters (x, y, z , width, length, height, heading) are converted into a vectorized BEV representation with N_d vertices, each modeled as:

$$\mathbf{v}_j^d = (\mu_x^{(j)}, b_x^{(j)}, \mu_y^{(j)}, b_y^{(j)}) \in \mathbb{R}^4, \quad j = 1, \dots, N_d.$$

Static and dynamic queries Q_s and Q_d are processed by probabilistic heads (MLPs) to generate vertex predictions:

$$\hat{V}_s = f_p(Q_s), \quad \hat{V}_d = f_p(Q_d), \quad (1)$$

where f_p predicts the Laplace distribution parameters for each vertex. The resulting vertex sets are given by

$$\hat{V}_s = \{\mathbf{v}_i^s\}_{i=1}^{N_s}, \quad \hat{V}_d = \{\mathbf{v}_j^d\}_{j=1}^{N_d}, \quad (2)$$

with each \mathbf{v}_i^s and \mathbf{v}_j^d parameterizing the location and scale in the BEV plane. To ensure strictly positive scale parameters and numerical stability, we apply a Softplus activation with a small offset ($\epsilon = 10^{-6}$) to the raw network outputs of scale parameter.

B. Uncertainty Fusion Module

The uncertainty fusion module encodes the predicted vertices into uncertainty features E_s^u and E_d^u via an MLP encoder $\text{Uncer}_{\text{enc}}$, given by

$$\begin{aligned} E_s^u &= \text{Uncer}_{\text{enc}}(\hat{V}_s) \in \mathbb{R}^{N_s \times d_h}, \\ E_d^u &= \text{Uncer}_{\text{enc}}(\hat{V}_d) \in \mathbb{R}^{N_d \times d_h}, \end{aligned} \quad (3)$$

where d_h is the hidden dimension, consistent with the query embedding dimension.

Given queries $Q_s \in \mathbb{R}^{N_s \times d_h}$ and $Q_d \in \mathbb{R}^{N_d \times d_h}$, we refine them by attending to the uncertainty features via multi-head cross-attention:

$$\begin{aligned} Q_s^{\text{uncer}} &= \text{Attention}_s(Q_s, E_s^u, E_s^u), \\ Q_d^{\text{uncer}} &= \text{Attention}_d(Q_d, E_d^u, E_d^u), \end{aligned} \quad (4)$$

producing uncertainty-aware queries Q_s^{uncer} and Q_d^{uncer} that replace Q_s and Q_d in downstream planning.

C. Uncertainty-aware Gate

The uncertainty-aware gate modulates historical information based on current scene uncertainty. Depending on the different AD approaches, it generates gating signals at different granularities: for high-dimensional temporal queries, it produces lightweight time-step-wise gating weights; for low-dimensional ego states, it generates a fine-grained feature-time matrix.

Let $Q^{\text{uncer}} \in \mathbb{R}^{N \times d_h}$ denote the current uncertainty-aware queries, where $N = N_s + N_d$. When applying the gate to temporal perception queries $Q_{\text{temporal}} \in \mathbb{R}^{N_{tq} \times T \times d_h}$ (where N_{tq} is large), it generates a vector $\mathbf{g}^{(t)} \in [0, 1]^T$ that weights each time step uniformly across all queries and features. First, the N uncertainty queries are aggregated into a global context $\mathbf{c} \in \mathbb{R}^{d_h}$ via attention:

$$\mathbf{c} = \sum_{i=1}^N \alpha_i \mathbf{q}_i, \quad \alpha_i = \frac{\exp(f_{\text{attn}}(\mathbf{q}_i))}{\sum_{j=1}^N \exp(f_{\text{attn}}(\mathbf{q}_j))}, \quad (5)$$

where \mathbf{q}_i is the i -th row of Q^{uncer} and $f_{\text{attn}}: \mathbb{R}^{d_h} \rightarrow \mathbb{R}$ is a linear layer. This context is then projected to T dimensions and passed through a sigmoid:

$$\mathbf{g}^{(t)} = \sigma(f_{\text{gate}}^{(t)}(\mathbf{c})), \quad (6)$$

with $f_{\text{gate}}^{(t)}: \mathbb{R}^{d_h} \rightarrow \mathbb{R}^T$. The gating vector is broadcast to shape $[1, T, 1]$ and applied element-wise:

$$Q_{\text{temporal}}^{\text{gated}} = \mathbf{g}^{(t)} \odot Q_{\text{temporal}}. \quad (7)$$

This lightweight design efficiently modulates all N_{tq} queries per time step using a single weight $g_t^{(t)}$. The modulated features $Q_{\text{temporal}}^{\text{gated}}$ are then passed to the original planner design, enabling dynamic, uncertainty-aware filtering of historical temporal information.

The same gating mechanism can be adapted to historical ego states $S_{\text{ego}} \in \mathbb{R}^{L \times T}$ (where L is the number of ego state features, e.g., turn signals, velocities, accelerations) by increasing the granularity of the gating signal. Instead of generating a vector $\mathbf{g}^{(t)} \in [0, 1]^T$, we produce a matrix $\mathbf{g}^{(e)} \in [0, 1]^{L \times T}$ using a separate projection layer $f_{\text{gate}}^{(e)}: \mathbb{R}^{d_h} \rightarrow \mathbb{R}^{L \times T}$ after the same context \mathbf{c} (Eq. (5)):

$$\mathbf{g}^{(e)} = \sigma(f_{\text{gate}}^{(e)}(\mathbf{c})). \quad (8)$$

This allows independent gating for each ego-state feature at every time step. The gated output is then computed as $S_{\text{ego}}^{\text{gated}} = \mathbf{g}^{(e)} \odot S_{\text{ego}}$, and replaces the original ego status input to the planner, keeping the rest of the model architecture unchanged.

IV. EXPERIMENTS

We evaluate UniUncer on nuScenes and NavsimV2, showing improved trajectory accuracy (7% lower L2 error) and driving score (10.8% higher EPDMS) over strong baselines. Ablation studies confirm that modeling uncertainty in both static and dynamic elements, combined with the uncertainty-aware gate, consistently enhances performance.

A. Main Results

Dataset. We use nuScenes [40] to evaluate our open-loop performance. The nuScenes dataset contains 1000 driving scenes recorded with six cameras and a 32-beam LiDAR sensor. It provides labels for dynamic objects, static map elements, and the trajectories of both the ego vehicle and dynamic objects, which are used for planning and prediction tasks. For pseudo closed-loop evaluation, we adopt NavsimV2 [44], a planning-oriented autonomous driving dataset built on top of OpenScene, which itself is a redistribution of nuPlan. Navsim [41] includes eight 1920×1080 cameras and a fused LiDAR point cloud. It was originally designed for open-loop evaluation of E2E methods. Later, NavsimV2 extended this work by rendering novel views with 3DGS to construct a pseudo closed-loop simulation.

Evaluation Metrics. For nuScenes, we report the L2 displacement error of the ego vehicle's future trajectory and the collision rate. For NavsimV2, we adopt the Extended Predictive Driver Model Score (EPDMS) [44] as the pseudo

TABLE II
PSEUDO CLOSE LOOP E2E PERFORMANCE BENCHMARK ON THE NAVHARD TWO-STAGE TEST FROM NAVSIMV2.

Method	Backbone	Stage	NC ↑	DAC ↑	DDC ↑	TLC ↑	EP ↑	TTC ↑	LK ↑	HC ↑	EC ↑	EPDMS Stage ↑	EPDMS ↑
LTF [39]	ResNet-34	Stage 1	96.2	79.6	99.1	99.6	84.1	95.1	94.2	97.6	79.1	62.3	23.1
	ResNet-34	Stage 2	77.8	70.4	84.3	98.1	85.1	75.7	45.4	95.7	76.0	37.5	
DiffusionDrive [16]	ResNet-34	Stage 1	96.3	84.7	98.7	99.3	84.3	95.3	94.4	97.8	76.4	65.5	25.9
	ResNet-34	Stage 2	79.5	71.2	84.9	98.3	86.5	75.6	46.9	96.2	70.9	39.6	
DiffusionDrive + Ours	ResNet-34	Stage 1	96.8	82.9	98.8	99.3	83.6	95.8	96.0	97.6	77.3	64.1	28.7
	ResNet-34	Stage 2	81.5	74.4	86.9	98.1	85.2	78.6	49.2	95.7	74.2	43.4	

closed-loop planning evaluation metric, particularly on its most challenging benchmark, the Navhard two-stage test. EPDMS is a weighted sum of several sub-metrics, including No At-Fault Collisions (NC), Drivable Area Compliance (DAC), Driving Direction Compliance (DDC), Lane Keeping (LK), Time-to-Collision (TTC), History Comfort (HC), Extended Comfort (EC), Traffic Light Compliance (TLC), and Ego Progress (EP), providing a comprehensive closed-loop planning score. The specific EPDMS is defined as:

$$\text{EPDMS} = \frac{\prod_{m \in \mathcal{M}_{\text{pen}}} \text{filter}_m(\text{agent}, \text{human}) \cdot \sum_{m \in \mathcal{M}_{\text{avg}}} w_m \text{filter}_m(\text{agent}, \text{human})}{\sum_{m \in \mathcal{M}_{\text{avg}}} w_m},$$

where $\mathcal{M}_{\text{pen}} = \{\text{NC}, \text{DAC}, \text{DDC}, \text{TLC}\}$ denotes the penalty terms, and $\mathcal{M}_{\text{avg}} = \{\text{TTC}, \text{EP}, \text{HC}, \text{LK}, \text{EC}\}$ denotes weighted average terms.

Evaluated Models. We implement our UniUncer on top of SparseDrive [15] for open-loop evaluation on nuScenes. We also incorporate UniUncer into DiffusionDrive [16] for comparative pseudo close-loop analysis. The main model architectures remain largely unchanged, with three key modifications: (1) replacing the deterministic heads for static and dynamic objects with probabilistic heads that estimate both μ and b associated with object vertices; (2) adding two *Uncertainty Fusion Modules* to incorporate the estimated uncertainties into static and dynamic queries; and (3) introducing an *Uncertainty-aware Gate* for gating temporal object-related features in SparseDrive for every time step. For DiffusionDrive, the *Uncertainty-aware Gate* is instead applied to the historical ego status, as it does not include temporal object features.

Evaluation on nuScenes. As shown in Tab.I, UniUncer improves the L2 trajectory displacement at all time steps, achieving a 7% reduction on average. Moreover, it achieves a lower collision rate, average collision rate reduced from 0.08% to 0.07%, compared to the baseline, demonstrating improved safety without sacrificing trajectory accuracy.

Evaluation on NavsimV2. As shown in Tab.II, our model improves the overall EPDMS by 10.8%, increasing from 25.9 to 28.7. The stage two EPDMS also improves significantly, from 39.6 to 43.4, while the stage one score experiences a slight decrease. We attribute this to the fact that stage two contains many synthetically generated challenging scenarios with artifacts that can confuse perception, where uncertainty estimation provides substantial benefits. This result high-

lights the effectiveness of our design in handling complex and safety-critical driving conditions where perception reliability is compromised.

Training Details. We keep the data format identical to the baselines, only modifying dynamic object labels by replacing box parameters with vertices. Four corner points and one center point serve as regression targets. The static and dynamic branches are supervised by a combination of ℓ_1 and Negative Log-Likelihood (NLL) losses:

$$L_{\text{static/dynamic}} = w_1 \cdot \ell_1^{\text{static/dynamic}} + w_2 \cdot \text{NLL}_{\text{static/dynamic}},$$

where w_1 and w_2 balance the losses ($w_1 = 0.25$, $w_2 = 0.6$ for dynamic; $w_1 = 0$, $w_2 = 1.0$ for static).

For nuScenes, we fine-tune from the official SparseDrive stage one checkpoint for 10 epochs to ensure convergence of the new probability heads, followed by another 10 epochs for stage two. For NavsimV2, we train from scratch using the same configuration as DiffusionDrive, since NavsimV2 excludes LiDAR data while the released DiffusionDrive weights were trained with LiDAR. All experiments use four NVIDIA RTX 3090 GPUs.

B. Ablation Study

We conduct ablation studies on both NavsimV2 and nuScenes to validate the contribution of each component in UniUncer.

Notation in the ablation tables. In Tab.III and Tab.IV, we use 'S-Uncer' and 'D-Uncer' to denote the uncertainty estimation for static and dynamic objects, respectively. 'Uncer-Gate' denotes the uncertainty-aware gate. Note that for the NavsimV2 ablation III, we focus on dynamic uncertainty and the gate, as the dataset does not provide vectorized map labels required for training the static uncertainty head. This highlights a practical limitation but also showcases the modular benefit of our framework and components can be applied independently where data permits.

Effectiveness of uncertainty estimation for dynamic objects and uncertainty-aware gate. As shown in Tab.III, we report the performance on the Navhard two-stage evaluation using DiffusionDrive as the baseline and progressively adding our proposed modules. With uncertainty estimation and the uncertainty fusion module, the EPDMS improves from 25.9 to 26.7. Incorporating all modules, including the uncertainty-aware gate, further increases the EPDMS to 28.7. These steady improvements demonstrate the effectiveness of our proposed modules, validating both the contribution of uncertainty estimation for dynamic objects and the role of

TABLE III

ABLATION STUDY ON THE EFFECTIVENESS OF DYNAMIC UNCERTAINTY AND THE UNCERTAINTY-AWARE GATE ON THE NAVHARD TWO-STAGE TEST.

Method	D-Uncer	Uncer-Gate	Stage	NC \uparrow	DAC \uparrow	DDC \uparrow	TLC \uparrow	EP \uparrow	TTC \uparrow	LK \uparrow	HC \uparrow	EC \uparrow	EPDMS Stage \uparrow	EPDMS \uparrow
DiffusionDrive [16]			Stage 1	96.3	84.7	98.7	99.3	84.3	95.3	94.4	97.8	76.4	65.4	25.9
			Stage 2	79.5	71.2	84.9	98.3	86.5	75.6	46.9	96.2	70.9	39.6	
+D-Uncer	\checkmark		Stage 1	95.9	85.6	98.7	99.3	84.5	95.1	94.9	97.8	76.0	65.5	26.7
			Stage 2	79.2	72.3	85.9	98.2	86.6	76.2	48.2	96.3	70.1	40.6	
+D-Uncer +Uncer-Gater	\checkmark	\checkmark	Stage 1	96.8	82.9	98.8	99.3	83.6	95.8	96.0	97.6	77.3	64.1	28.7
			Stage 2	81.5	74.4	86.9	98.9	85.2	78.6	49.2	95.7	74.2	43.4	

TABLE IV

ABLATION STUDY ON THE EFFECTIVENESS OF STATIC UNCERTAINTY, DYNAMIC UNCERTAINTY, AND THE UNCERTAINTY-AWARE GATE ON NUSCENES.

Method	S-Uncer	D-Uncer	Uncer-Gate	L2 (m) \downarrow				Col. Rate (%) \downarrow			
				1s	2s	3s	Avg.	1s	2s	3s	Avg.
SparseDrive [15]				0.29	0.58	0.96	0.61	0.01	0.05	0.18	0.08
+S-Uncer	\checkmark			0.30	0.58	0.95	0.61	0.01	0.05	0.18	0.08
+D-Uncer		\checkmark		0.27	0.53	0.90	0.57	0.02	0.11	0.21	0.11
+S-Uncer +D-Uncer	\checkmark	\checkmark		0.27	0.53	0.88	0.56	0.01	0.06	0.19	0.09
+S-Uncer +D-Uncer +Uncer-Gate	\checkmark	\checkmark	\checkmark	0.27	0.53	0.90	0.57	0.00	0.04	0.17	0.07

uncertainty-aware gate for past ego status. We do not involve static uncertainty, since the vectorized label for the map element is not available in NavsimV2.

Effectiveness of uncertainty estimation for dynamic and static objects and uncertainty-aware gate. Tab.IV presents an ablation study in which we gradually incorporate our proposed modules into SparseDrive on nuScenes data set. Introducing uncertainty for static objects alone provides limited improvement in L2 displacement, consistent with results reported in UncAD [17]. When uncertainty is applied to dynamic objects only, the average L2 displacement error decreases from 0.61 m to 0.57 m, although the collision rate slightly increases. Incorporating both static and dynamic uncertainty leads to improvements in both L2 error and collision rate, reducing the average L2 error to 0.56 m while the average collision rate only slightly increases from 0.08% to 0.09%. Finally, when the full model, combining static and dynamic uncertainty estimation with the uncertainty-aware gate, we observe consistent gains in both metrics: the average L2 error decreases to 0.57 m, and the average collision rate drops to 0.07%. These results demonstrate that explicitly modeling both static and dynamic uncertainties, together with adaptive gating of historical features, effectively enhances trajectory prediction accuracy while reducing collisions, highlighting the practical benefit of our unified uncertainty framework for E2E autonomous driving.

Efficiency of Our Model. We compare our model with SparseDrive, which serves as the backbone of our method. As shown in Tab.V, our model introduces only 5.6 M additional parameters compared to the baseline, while incurring just a 0.5 FPS reduction. This result aligns with our design goal: enabling unified uncertainty estimation with only minimal modifications to the original architecture.

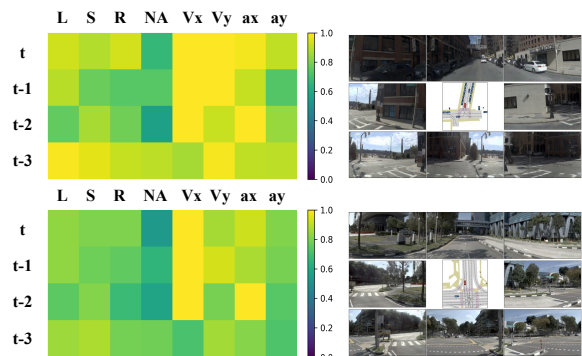


Fig. 3. Visualization of the Uncertainty-aware Gate (fine-grained feature-time matrix mode) for Ego Status. The upper row illustrates a complex scenario, while the lower row corresponds to a simple scenario. Driving commands include *L* (left turn), *S* (keep straight), *R* (right turn), and *NA* (not available). V_x and V_y denote velocities along the x - and y -axes, respectively, while a_x and a_y represent accelerations along the x - and y -axes.

TABLE V

EFFICIENCY OF OUR APPROACH. ALL METHODS ARE TESTED ON A SINGLE RTX 3090 GPU.

Method	Number of Params (M)	FPS
SparseDrive [15]	86.2	6.6
SparseDrive + Ours	91.8	6.1

C. Qualitative Results

We visualize model behavior to demonstrate the effectiveness of our modules.

Using ego status from t to $t - 3$ (commands, velocities, accelerations), Fig. 3 shows that the uncertainty-aware gate emphasizes richer historical information in complex multi-object scenes, while prioritizing x -axis velocity and acceleration in simpler scenes.

Fig. 4 displays estimated uncertainties and planned trajectories on nuScenes. Well-calibrated x - and y -axis uncertain-

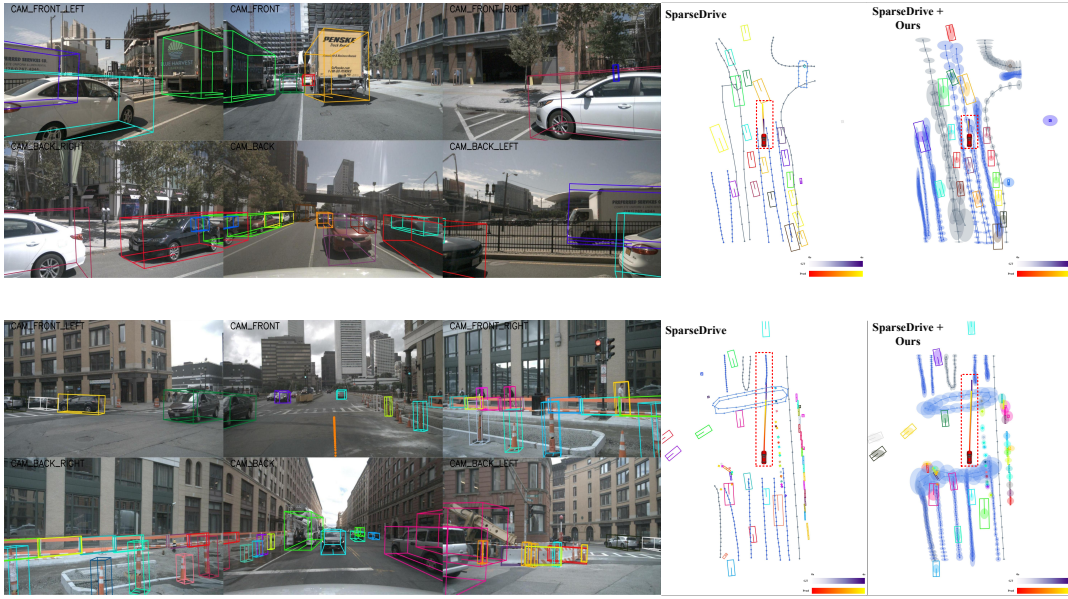


Fig. 4. Qualitative results on the nuScenes dataset. With unified uncertainty estimation, the predicted trajectory from our model is closer to the ground-truth trajectory compared to the baseline model. Ellipses on map vertices and dynamic objects represent the scale parameter b along the x - and y -axes. For dynamic objects, the uncertainty is lower on the visible side of the vehicle compared to the invisible side. Best viewed in the digital version.

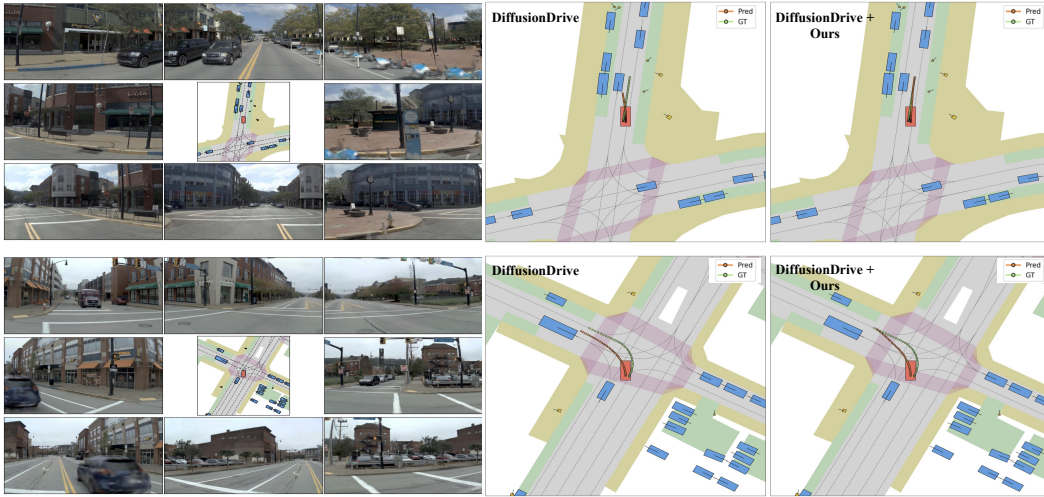


Fig. 5. Qualitative results on the Navhard two-stage test. Uncertainty estimation enables the planner to generate trajectories that are closer to the ground-truth (GT) trajectory in general. It demonstrates that uncertainty estimation for dynamic agents is critical for collision avoidance in challenging pseudo closed-loop scenarios rendered by 3DGS, which can introduce synthetic artifacts (e.g., the bus and vehicle in the second row). By assigning higher uncertainty to such potentially unreliable detections, the planner adopts a more conservative strategy, successfully avoiding over-reaction and improving robustness. Better view in digital version.

ties help the planner produce trajectories closer to human demonstrations, indicating more natural and robust behavior. (Uncertainties per vehicle are aggregated at its center for visualization.)

Fig. 5 illustrates that dynamic uncertainty is critical for collision avoidance in challenging 3DGS-rendered scenes, where synthetic objects may exhibit artifacts. By assigning higher uncertainty to unreliable detections, the planner adopts a conservative strategy, avoiding over-reaction and improving safety.

V. CONCLUSION

We present UniUncer, a lightweight and plug-and-play framework that models uncertainty in both dynamic and

static objects for end-to-end driving. It leverages this unified uncertainty to adaptively gate historical information such as ego states and temporal queries. Experiments on nuScenes and NavsimV2 show consistent improvements in open-loop and pseudo closed-loop evaluations, enhancing both trajectory accuracy and safety. The design adds minimal overhead and is readily adaptable to various E2E backbones. This work encourages future research on interactive uncertainty and epistemic uncertainty estimation for out-of-distribution robustness.

REFERENCES

- [1] S. Hu, L. Chen, P. Wu, H. Li, J. Yan, and D. Tao, "St-p3: End-to-end vision-based autonomous driving via spatial-temporal feature learning," in *ECCV*, 2022, pp. 533–549.
- [2] 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," *NeurIPS*, vol. 35, pp. 6119–6132, 2022.
- [3] W. Tong, C. Sima, T. Wang, L. Chen, S. Wu, H. Deng, Y. Gu, L. Lu, P. Luo, D. Lin, *et al.*, "Scene as occupancy," in *ICCV*, 2023.
- [4] 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 *CVPR*, 2023, pp. 21 983–21 994.
- [5] 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 *CVPR*, 2023, pp. 17 853–17 862.
- [6] 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 *ICCV*, 2023, pp. 8340–8350.
- [7] Z. Yan, P. Li, Z. Fu, S. Xu, Y. Shi, X. Chen, Y. Zheng, Y. Li, T. Liu, C. Li, *et al.*, "Int2: Interactive trajectory prediction at intersections," in *ICCV*, 2023, pp. 8536–8547.
- [8] B. Tian, M. Liu, H.-a. Gao, P. Li, H. Zhao, and G. Zhou, "Unsupervised road anomaly detection with language anchors," in *ICRA*, 2023, pp. 7778–7785.
- [9] X. Jia, Y. Gao, L. Chen, J. Yan, P. L. Liu, and H. Li, "Driveadapter: Breaking the coupling barrier of perception and planning in end-to-end autonomous driving," in *ICCV*, 2023, pp. 7953–7963.
- [10] B. Jin, Y. Zheng, P. Li, W. Li, Y. Zheng, S. Hu, X. Liu, J. Zhu, Z. Yan, H. Sun, *et al.*, "Tod3cap: Towards 3d dense captioning in outdoor scenes," in *ECCV*, 2024, pp. 367–384.
- [11] 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 *CVPR*, 2024, pp. 14 864–14 873.
- [12] Z. Chen, M. Ye, S. Xu, T. Cao, and Q. Chen, "Ppad: Iterative interactions of prediction and planning for end-to-end autonomous driving," in *ECCV*, 2024, pp. 239–256.
- [13] W. Zheng, R. Song, X. Guo, C. Zhang, and L. Chen, "Genad: Generative end-to-end autonomous driving," in *ECCV*, 2024.
- [14] X. Weng, B. Ivanovic, Y. Wang, Y. Wang, and M. Pavone, "Para-drive: Parallelized architecture for real-time autonomous driving," in *CVPR*, 2024, pp. 15 449–15 458.
- [15] W. Sun, X. Lin, Y. Shi, C. Zhang, H. Wu, and S. Zheng, "Sparsedrive: End-to-end autonomous driving via sparse scene representation," in *ICRA*, 2025, pp. 8795–8801.
- [16] 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," in *CVPR*, 2025.
- [17] P. Yang, Y. Zheng, Q. Zhang, K. Zhu, Z. Xing, Q. Lin, Y.-F. Liu, Z. Su, and D. Zhao, "Uncad: Towards safe end-to-end autonomous driving via online map uncertainty," in *ICRA*, 2025, pp. 6408–6415.
- [18] 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," in *CVPR*, 2025, pp. 1602–1611.
- [19] B. Zhang, N. Song, X. Jin, and L. Zhang, "Bridging past and future: End-to-end autonomous driving with historical prediction and planning," in *CVPR*, 2025, pp. 6854–6863.
- [20] Z. Song, C. Jia, L. Liu, H. Pan, Y. Zhang, J. Wang, X. Zhang, S. Xu, L. Yang, and Y. Luo, "Don't shake the wheel: Momentum-aware planning in end-to-end autonomous driving," in *CVPR*, 2025.
- [21] Y. Li, L. Fan, J. He, Y. Wang, Y. Chen, Z. Zhang, and T. Tan, "Enhancing end-to-end autonomous driving with latent world model," *ICLR*, 2025.
- [22] P. Li and D. Cui, "Navigation-guided sparse scene representation for end-to-end autonomous driving," *ICLR*, 2025.
- [23] X. Jia, J. You, Z. Zhang, and J. Yan, "Drivetransformer: Unified transformer for scalable end-to-end autonomous driving," *ICLR*, 2025.
- [24] C. Li, K. Zhou, T. Liu, Y. Wang, M. Zhuang, H.-a. Gao, B. Jin, and H. Zhao, "Avd2: Accident video diffusion for accident video description," in *ICRA*, 2025, pp. 13 289–13 296.
- [25] Z. Zhang, X. Li, S. Zou, G. Chi, S. Li, X. Qiu, G. Wang, G. Zheng, L. Wang, H. Zhao, *et al.*, "Chameleon: Fast-slow neuro-symbolic lane topology extraction," *arXiv preprint arXiv:2503.07485*, 2025.
- [26] K. Ding, B. Chen, Y. Su, H.-a. Gao, B. Jin, C. Sima, X. Li, W. Zhang, P. Barsch, H. Li, *et al.*, "Hint-ad: Holistically aligned interpretability in end-to-end autonomous driving," in *CoRL*, 2025, pp. 3742–3765.
- [27] Y. Li, Z. Zhang, X. Qiu, X. Li, Z. Liu, L. Wang, R. Li, Z. Zhu, H.-a. Gao, X. Lin, *et al.*, "Reusing attention for one-stage lane topology understanding," in *IROS*, 2025, pp. 16977–16984.
- [28] H. Chi, H.-a. Gao, Z. Liu, J. Liu, C. Liu, J. Li, K. Yang, Y. Yu, Z. Wang, W. Li, *et al.*, "Impromptu via: Open weights and open data for driving vision-language-action models," *NeurIPS*, 2025.
- [29] Y. Zheng, P. Yang, Z. Xing, Q. Zhang, Y. Zheng, Y. Gao, P. Li, T. Zhang, Z. Xia, P. Jia, *et al.*, "World4drive: End-to-end autonomous driving via intention-aware physical latent world model," in *ICCV*, 2025, pp. 28 632–28 642.
- [30] Y. Lu, J. Tu, Y. Ma, and X. Zhu, "Real-ad: Towards human-like reasoning in end-to-end autonomous driving," in *ICCV*, 2025.
- [31] H. Fu, D. Zhang, Z. Zhao, J. Cui, D. Liang, C. Zhang, D. Zhang, H. Xie, B. Wang, and X. Bai, "Orion: A holistic end-to-end autonomous driving framework by vision-language instructed action generation," in *ICCV*, 2025, pp. 24 823–24 834.
- [32] Y. Tang, Z. Xu, Z. Meng, and E. Cheng, "Hip-ad: Hierarchical and multi-granularity planning with deformable attention for autonomous driving in a single decoder," in *ICCV*, 2025, pp. 25 605–25 615.
- [33] J. Guo, Y. Ding, X. Chen, S. Chen, B. Li, Y. Zou, X. Lyu, F. Tan, X. Qi, Z. Li, *et al.*, "Dist-4d: Disentangled spatiotemporal diffusion with metric depth for 4d driving scene generation," in *ICCV*, 2025.
- [34] 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.
- [35] A. Jiang, Y. Gao, Z. Sun, Y. Wang, J. Wang, J. Chai, Q. Cao, Y. Heng, H. Jiang, Y. Dong, *et al.*, "Diffvla: Vision-language guided diffusion planning for autonomous driving," *arXiv preprint arXiv:2505.19381*, 2025.
- [36] Z. Zhang, X. Qiu, B. Zhang, G. Zheng, X. Gu, G. Chi, H.-a. Gao, L. Wang, Z. Liu, X. Li, *et al.*, "Delving into mapping uncertainty for mapless trajectory prediction," in *IROS*, 2025, pp. 16 969–16 976.
- [37] Y. Li, Y. Wang, Y. Liu, J. He, L. Fan, and Z. Zhang, "End-to-end driving with online trajectory evaluation via bev world model," in *ICCV*, 2025, pp. 27 137–27 146.
- [38] X. Gu, G. Song, I. Gilitschenski, M. Pavone, and B. Ivanovic, "Producing and leveraging online map uncertainty in trajectory prediction," in *CVPR*, 2024, pp. 14 521–14 530.
- [39] 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.
- [40] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusscenes: A multimodal dataset for autonomous driving," in *CVPR*, 2020.
- [41] 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," *NeurIPS*, vol. 37, pp. 28 706–28 719, 2024.
- [42] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *CoRL*, 2017, pp. 1–16.
- [43] X. Jia, Z. Yang, Q. Li, Z. Zhang, and J. Yan, "Bench2drive: Towards multi-ability benchmarking of closed-loop end-to-end autonomous driving," *NeurIPS*, vol. 37, pp. 819–844, 2024.
- [44] W. Cao, M. Hallgarten, T. Li, D. Dauner, X. Gu, C. Wang, Y. Miron, M. Aiello, H. Li, I. Gilitschenski, *et al.*, "Pseudo-simulation for autonomous driving," in *CoRL*, 2025, pp. 4709–4722.
- [45] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3d gaussian splatting for real-time radiance field rendering," *ACM Trans. Graph.*, vol. 42, no. 4, pp. 139–1, 2023.
- [46] A. Kendall and Y. Gal, "What uncertainties do we need in bayesian deep learning for computer vision?" *NeurIPS*, vol. 30, 2017.
- [47] J. Gast and S. Roth, "Lightweight probabilistic deep networks," in *CVPR*, 2018, pp. 3369–3378.
- [48] M. Poggi, F. Aleotti, F. Tosi, and S. Mattoccia, "On the uncertainty of self-supervised monocular depth estimation," in *CVPR*, 2020.
- [49] G. P. Meyer, A. Laddha, E. Kee, C. Vallespi-Gonzalez, and C. K. Wellington, "Lasernet: An efficient probabilistic 3d object detector for autonomous driving," in *CVPR*, 2019, pp. 12 677–12 686.