

CoPlanner: An Interactive Motion Planner with Contingency-Aware Diffusion for Autonomous Driving

Ruiguo Zhong, Ruoyu Yao, Pei Liu, Xiaolong Chen, Rui Yang, Jun Ma, *Senior Member, IEEE*

Abstract—Accurate trajectory prediction and motion planning are crucial for autonomous driving systems to navigate safely in complex, interactive environments characterized by multimodal uncertainties. However, current generation-then-evaluation frameworks typically construct multiple plausible trajectory hypotheses but ultimately adopt a single most likely outcome, leading to overconfident decisions and a lack of fallback strategies that are vital for safety in rare but critical scenarios. Moreover, the usual decoupling of prediction and planning modules could result in socially inconsistent or unrealistic joint trajectories, especially in highly interactive traffic. To address these challenges, we propose a contingency-aware diffusion planner (CoPlanner), a unified framework that jointly models multi-agent interactive trajectory generation and contingency-aware motion planning. Specifically, the pivot-conditioned diffusion mechanism anchors trajectory sampling on a validated, shared short-term segment to preserve temporal consistency, while stochastically generating diverse long-horizon branches that capture multimodal motion evolutions. In parallel, we design a contingency-aware multi-scenario scoring strategy that evaluates candidate ego trajectories across multiple plausible long-horizon evolution scenarios, balancing safety, progress, and comfort. This integrated design preserves feasible fallback options and enhances robustness under uncertainty, leading to more realistic interaction-aware planning. Extensive closed-loop experiments on the nuPlan benchmark demonstrate that CoPlanner consistently surpasses state-of-the-art methods on both Val14 and Test14 datasets, achieving significant improvements in safety and comfort under both reactive and non-reactive settings. The source code is available on GitHub¹.

I. INTRODUCTION

Trajectory prediction and motion planning are fundamental components of autonomous driving, which enable the anticipation of surrounding agents' behaviors and facilitate informed decision-making to ensure safety and efficiency in complex driving scenarios [1], [2]. Recent advances in scene-level multi-agent trajectory generation, which models the joint distribution of all agents' trajectories, demonstrate that diffusion models generate diverse, multimodal futures [3], [4]. This generation step is typically followed by an evaluator that selects an executable plan, forming a generation-then-evaluation framework [5], [6]. However, collapsing to a single most-likely future leaves the planner unprepared for alternative evolutions. When the scene unfolds differently, it cannot accommodate the broader set of plausible outcomes,

Ruiguo Zhong, Ruoyu Yao, Pei Liu, Rui Yang, Xiaolong Chen, and Jun Ma are with The Hong Kong University of Science and Technology (Guangzhou), Guangzhou 511453, China (e-mail: {rzhong151, ryao092, pliu061, ryang253}@connect.hkust-gz.edu.cn; xiaolongchen@hkust-gz.edu.cn; jun.ma@ust.hk). (*Corresponding Author: Jun Ma.*)

¹<https://github.com/ruiguoz/Contingency-Diffusion-Planner>

leading to late braking, unstable execution, or unsafe maneuvers.

Contingency or risk-aware planning mitigates some of these issues by maintaining multiple candidate plans that share a short-term segment and branch later, deferring commitment until uncertainty resolves [7], [8], [9]. However, in optimization-based formulations, predictions are usually treated as exogenous inputs and the planner optimizes trajectories conditioned on them. This decouples prediction from planning [10], [11], [12], which yields unrealistic or socially inconsistent joint trajectories, particularly in highly interactive or congested situations.

Recent generation-then-evaluation frameworks sample multiple futures and rank them with an evaluator [5], [6]. In practice, the evaluator scores each scenario independently and then collapses to a single highest-scored one, discarding valuable fallback options and ignoring the mutual exclusivity of alternative behaviors for the same agent. In addition, without a shared short-term anchor, candidates often diverge too early, which harms execution stability and yields plans that are inconsistent under equally plausible evolutions. This mismatch makes it challenging for learning-based planners to naturally integrate contingency planning into their decision-making process.

Robust planning requires both preserving a consistent short-term execution and retaining feasible long-term trajectories for contingency. To this end, we propose a **Contingency-Aware Diffusion Planner (CoPlanner)**, a unified generation-then-evaluation framework that explicitly considers contingency with diffusion-based motion planning. A key innovation of our method is a pivot-conditioned diffusion mechanism, which anchors all sampled joint trajectories to a validated shared short-term segment while generating diverse long-term branches to capture multimodal futures. To align short-term stability with long-term preparedness, we introduce a contingency-aware multi-scenario scoring rule that evaluates each ego candidate plan across multiple generated futures, balancing safety, progress, and comfort. By coupling multimodal, interaction-aware generation with short-term segment consistency and multi-scenario evaluation, CoPlanner maintains feasible backup options and improves decision-making when uncertainty matters most. Our contributions are summarized as follows:

- We introduce CoPlanner, a unified generation-then-evaluation framework that jointly samples ego plans and also the trajectories of other agents, preserving near-term consistency while keeping contingency branches. This ensures the ego vehicle remains feasible across plausible

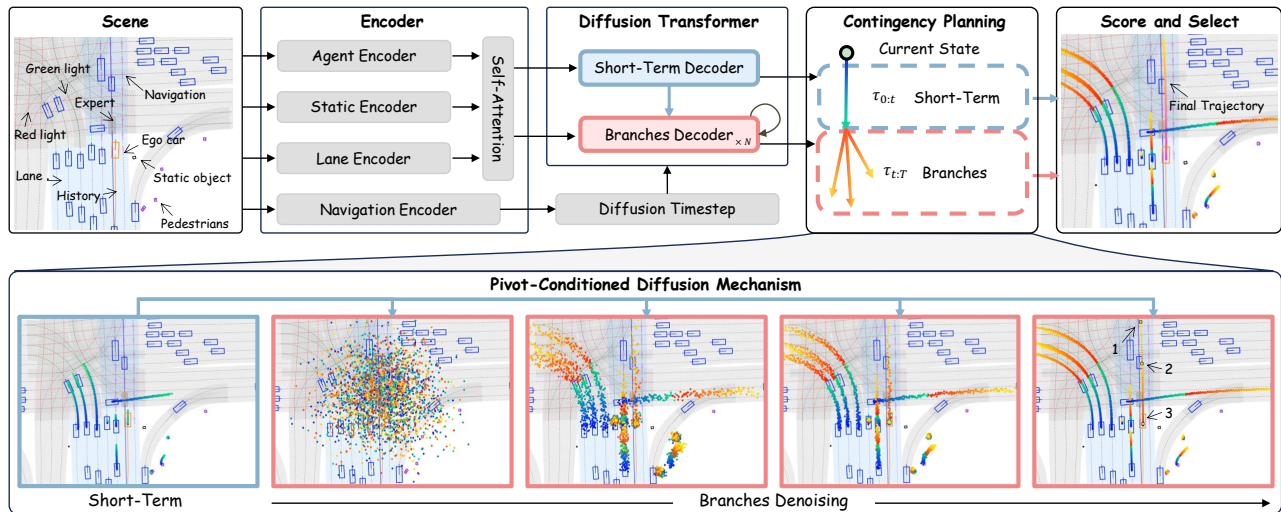


Fig. 1. Overview of CoPlanner. From vectorized scene context, the planner first generates a set of shared short-term segments for all agents. Conditioned on these anchors, pivot-conditioned diffusion mechanism completes multiple diverse long-horizon branches via masked reverse denoising. A contingency-aware scorer evaluates across scenarios and selects the final ego trajectory.

futures and enhances overall driving performance.

- We propose a pivot-conditioned diffusion mechanism, a masked reverse-diffusion scheme that freezes a validated shared short-term segment and generates diverse long-term branches. This improves temporal consistency while preserving multimodality.
- We develop a contingency-aware scoring strategy that pre-filters candidates on a short-term horizon and then aggregates branch scores across multiple generated futures. This mitigates worst-case over-conservatism and selects plans that remain feasible, safe, and comfortable across scenarios, without sacrificing driving efficiency.
- We conduct comprehensive closed-loop evaluations on nuPlan, achieving state-of-the-art performance on Val14 and Test14, with improved safety and comfort metrics.

II. RELATED WORK

A. Multi-Modal Behavior Prediction

Forecasting multiple plausible futures is essential for safe planning in interactive traffic. Vectorized scene representations that encode lanes, topology, and agent histories as polylines markedly improve the fidelity of learned predictors [13], [14]. Prior multimodal methods emphasize per-agent diversity [15], [16], but often struggle to produce coherent scene-level joint behaviors because inter-agent interactions are weakly modeled. Generative approaches address this by sampling plausible joint futures one at a time, thereby capturing both multimodality and cross-agent dependencies. Diffusion-based predictors, in particular, demonstrate strong diversity and realism in trajectory generation [3]. However, most such models remain downstream-agnostic: they generate agent futures without considering how a planner will exploit or contradict them. This disconnect can yield ego plans that appear feasible in isolation but are socially inconsistent once other agents react. Our work targets this gap by pairing multimodal generation with a mechanism that enforces a shared, safety-validated short-term segment across all candidates, making the

predictions directly usable by a contingency-aware planner.

B. Integrated Predictions and Planning

Traditional autonomous driving stacks separate prediction and planning: forecasts are produced first and then treated as fixed inputs to the planner. This modularity simplifies engineering but can under-represent feedback between the ego vehicle and other agents in interactive scenes. To capture such coupling, recent research explores integrated formulations that allow mutual influence between prediction and planning [1], [17]. One strand models interaction with game theory, which iteratively updates each agent’s policy via best-response or equilibrium computations. Recent work addresses intent uncertainty and improves scalability by casting interactions as potential games or trajectory-tree-based pairwise games [18], [19]. These methods are sensitive to intent priors and assumptions about other agents. A related line learns interaction patterns from data within a level- k framework, including GameFormer [20] and HPP [21], while still relying on iterative rollouts. In practice, such iterative trajectory generation can yield jittery closed-loop behavior. Beyond these, learning-based approaches treat interactions implicitly. Specifically, imitation-style planners like CALMMDrive [22] generate diverse ego candidates and select the highest-scoring one under a vision-language metric. Other approaches, such as Int2Planner [23], add an explicit prediction head and score trajectories jointly, yet typically execute only the single top-confidence trajectory. Generative scene roll-outs, notably GenDrive [6], further produce joint trajectories of ego vehicle and other agents and rank them with a vision-language model. In practice, many integrated systems still collapse to a single most-likely outcome at decision time. This discards valuable fallback options, overlooks the mutual exclusivity of alternative behaviors for the same agent, and can cause early divergence among candidates, harming short-term executability. The approach proposed in this paper complements these lines by explicitly enforcing near-term consistency (a shared

short-term segment) while retaining long-term diversity for contingency.

C. Contingency Planning

Unlike prior work that either focuses on multi-modal prediction, or contingency-aware planning without joint modeling, our method addresses the problem of consistent execution across scenarios by introducing a unified framework that jointly models a shared short-term segment and diverse long-term branches across futures [24], [25]. CMPC [9] and RACP [8] synthesize a bundle of trajectories that share an initial safe segment before branching, thereby hedging against model error or perception noise. Sampling-based approaches, such as LookOut [7], augment this idea with stochastic rollout and collision-set pruning. While effective in low-interaction scenes, these planners presume that every other agent follows a fixed, independently generated forecast. Moreover, many contingency planners implicitly marginalize over futures, requiring the ego vehicle to trade low-probability but high-impact events against nominal progress without explicitly reasoning about mutually exclusive outcomes. To explicitly model evolving interactions, tree-structured methods like TPP [26], DTPP [27] employ recursive generation to construct scenario trees. In contrast to such recursive architectures, diffusion models demonstrate far superior capability in covering complex multimodal distributions without suffering from compounding errors. Our framework couples contingency planning with scene-level joint trajectories generation for the ego vehicle and other agents by enforcing a shared short-term segment and sampling diverse long-term branches, which ensures feasibility across multiple plausible evolutions.

In summary, prior works suffer from three main limitations: (i) multimodal predictors generate diverse, coherent joint futures but are decoupled from planning, with no mechanism for using those samples in decision making; (ii) integrated predictions and planning systems often collapse to a single most-likely hypothesis at decision time, leaving no fallback when reality deviates and causing policy chattering; and (iii) contingency planners keep a shared prefix with late branching but typically optimize against exogenous, fixed forecasts, which weakens social consistency. We bridge these gaps by combining short-term conditioned diffusion with multi-scenario evaluation, yielding plans that are both stable to execute and robust under uncertainty.

III. PROBLEM FORMULATION

We consider a driving scene with M agents indexed by $j \in \{0, 1, \dots, M-1\}$, where $j=0$ denotes the ego vehicle and $j \geq 1$ denote the other agents. Time is discretized as $\{0, 1, \dots, T\}$ with step Δt . The future trajectory of agent j is $\tau^j = \{x_t^j\}_{t=0}^T$, and the joint future is $\Lambda = \{\tau^j\}_{j=0}^{M-1}$. Context \mathcal{C} includes agent histories, HD map, traffic signals, and route.

The classical objective of autonomous driving is to seek one ego plan $\tau^0 \in \mathcal{T}$ minimizing a nominal cost J under feasibility constraints:

$$\tau^* = \arg \min_{\tau \in \mathcal{T}} J(\tau) \quad (1)$$

where \mathcal{T} is the set of feasible candidate trajectories given the current ego state.

This single-scenario view implicitly commits to one most-likely evolution of the environment. To explicitly consider contingency, we introduce a branching time $t_b \in \{0, 1, \dots, T\}$ and split the joint future into a shared short-term segment $S = \{\tau_{0:t_b}^j\}_{j=0}^{M-1}$ and a long-term branch $B = \{\tau_{t_b:T}^j\}_{j=0}^{M-1}$ so that $\Lambda = [S, B]$. For an ego trajectory τ^0 and some potential future Λ , we decompose the evaluation window:

$$J(\tau^0 \mid \Lambda, \mathcal{C}) = J_{\text{shared}}(\tau_{0:t_b}^0 \mid \Lambda, \mathcal{C}) + J_{\text{branch}}(\tau_{t_b:T}^0 \mid \Lambda, \mathcal{C}). \quad (2)$$

A contingency policy executes an identical short-term ego trajectory segment under all plausible evolutions and branches later. Denote the ego anchor by $\hat{\tau}_{0:t_b}^0$ and the feasible set with that anchor by

$$\mathcal{T}_{\text{short}}(\hat{\tau}_{0:t_b}^0) = \{\tau^0 \in \mathcal{T} : \tau_{0:t_b}^0 = \hat{\tau}_{0:t_b}^0\}. \quad (3)$$

Rather than treating Λ as exogenous [8], we model the joint future generatively by factorizing it into a shared short-term segment S and a long-term branch B :

$$p(\Lambda \mid \mathcal{C}) = p_{\theta}(S \mid \mathcal{C}) p_{\phi}(B \mid S, \mathcal{C}). \quad (4)$$

Here, θ and ϕ parameterize the short-term generator and the conditional long-term completion model, respectively. At test time we sample K short-term segments and, for each segment, sample N long-term branches:

$$\hat{S}_k \sim p_{\theta}(S \mid \mathcal{C}), B_{k,n} \sim p_{\phi}(B \mid \hat{S}_k, \mathcal{C}), \Lambda_{k,n} = [\hat{S}_k, B_{k,n}], \quad (5)$$

with $k = 1, 2, \dots, K$ and $n = 1, 2, \dots, N$.

For a retained short-term segment \hat{S}_k we select one ego plan that shares $\hat{\tau}_{0:t_b}^{0,k}$ and is robust across its N plausible futures:

$$J_k(\tau^0) = J_{\text{shared}}(\tau_{0:t_b}^0 \mid \hat{S}_k, \mathcal{C}) \quad (6)$$

$$+ \mathcal{R}\left(\{J_{\text{branch}}(\tau_{t_b:T}^0 \mid \Lambda_{k,n}, \mathcal{C})\}_{n=1}^N\right),$$

$$\tilde{\tau}^{0,k} = \arg \min_{\tau^0 \in \mathcal{T}_{\text{short}}(\hat{\tau}_{0:t_b}^{0,k})} J_k(\tau^0). \quad (7)$$

In this work, we instantiate the risk aggregator \mathcal{R} as the equal-weight mean over the generated futures $\Lambda_{k,n}$.

IV. METHODOLOGY

A. Overview

An overview of the full pipeline is shown in Fig. 1. CoPlanner executes a generation-then-evaluation loop at each cycle: (i) **Short-term segment generation** produces K shared short-term segments $\{\hat{S}_k\}$ and keeps those passing hard checks on $[0, t_b]$; (ii) for each retained short-term segment, **pivot-conditioned diffusion mechanism** samples N diverse long-term branches $\{B_{k,n}\}$ to form $\{\Lambda_{k,n}\}$ that all share the same near-term joint evolution; (iii) **two-stage contingency scoring** first pre-filters ego candidates on $[0, t_b]$, then aggregates branch costs over $\{\Lambda_{k,n}\}$ on $(t_b, T]$ to obtain $\tilde{\tau}^{0,k}$; (iv) a selection across short-term segment groups outputs the final $\tau^{0,*}$. This design separates near-term stability from long-term preparedness while keeping them consistent.

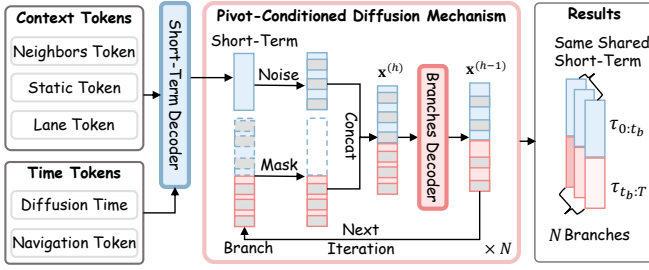


Fig. 2. Pivot-conditioned diffusion mechanism with masked reverse diffusion. At each reverse step, short-term frames ($t \leq t_b$) are re-noised from the fixed anchor and overwrite the state, while long-term frames ($t > t_b$) are denoised by G_{full} and merged via the mask. This guarantees a shared near-term execution and late branching.

B. Diffusion-based Scenario Generator

Backbone and Objective. We adopt a Diffusion Transformer (DiT) backbone operating on stacked scene trajectories for the ego vehicle and other agents. We use a variance-preserving (VP) diffusion process with continuous time $h \in [0, 1]$ that is discretized into H steps at inference. At each diffusion step h , the noised future $\mathbf{x}^{(h)}$ is fused with the scene context \mathbf{C} via gated cross-attention and a sinusoidal timestep embedding. To improve stability under heavy conditioning, we use the $\mathbf{x}^{(0)}$ -prediction objective, training the denoiser to directly reconstruct the clean sample:

$$\mathcal{L}_\theta = \mathbb{E}_{\mathbf{x}^{(0)}, h \sim \mathcal{U}(0,1), \mathbf{x}^{(h)} \sim q_{h|0}} \left[\|\mu_\theta(\mathbf{x}^{(h)}, h, \mathbf{C}) - \mathbf{x}^{(0)}\|_2^2 \right]. \quad (8)$$

For sampling, the continuous VP schedule is discretized to $\{\alpha_h\}_{h=1}^H$ with $\bar{\alpha}_h = \prod_{s=1}^h \alpha_s$, which matches the continuous schedule at the corresponding time, ensuring train–test consistency.

Short-Term Model G_{shared} . G_{shared} models the distribution of shared short-horizon joint futures $p_\theta(S | \mathbf{C})$ and is trained to produce high-confidence short-term segments over the interval $[0, t_b]$. The input is the scene context \mathbf{C} containing agent histories, a vectorized map, traffic controls, and the route. It uses the same diffusion backbone and the $\mathbf{x}^{(0)}$ -prediction objective as the full model, with the loss applied only on the short-term window.

Full Model G_{full} . G_{full} models the conditional distribution of long-term branches, $p_\phi(B | S, \mathbf{C})$, and completes the scene evolution given a fixed shared short-term segment. Its inputs are the scene context \mathbf{C} and a clean anchor short-term segment S , and its output is the joint trajectory on $[0, T]$. Training only on the long-term portion deprives the model of exposure to early-phase dynamics and interactions, which leads to mode collapse, boundary artifacts around t_b , and violations of vehicle kinematics. To avoid this, we predict the full horizon while applying a long-term corruption mask that injects noise only for $t > t_b$. The $\mathbf{x}^{(0)}$ -prediction objective is defined on $[0, T]$, whereas the diffusion loss is evaluated on branch frames. The clean short-term segment is provided as a conditioning signal, with optional short-term conditioning dropout to prevent over-reliance.

Algorithm 1 Pivot-Conditioned Diffusion Mechanism

```

1: Input: Context  $\mathbf{C}$ , Short-Term Segment  $\hat{S}$ , Mask  $m(t) = \mathbb{1}[t \leq t_b]$ , steps  $H$ 
2: for  $n = 1$  to  $N$  do ▷ Sample  $N$  branches
3:    $\mathbf{x}^{(H)} \sim \mathcal{N}(0, I)$ 
4:   for  $h = H$  to  $1$  do
5:      $z \sim \mathcal{N}(0, I)$  if  $h > 1$  else  $z = 0$ 
6:      $\mathbf{x}_{\text{short}}^{(h-1)} \leftarrow \text{re-noise}(\hat{S})$  ▷ Freeze short-term segment
7:      $\mathbf{x}_{\text{branch}}^{(h-1)} \leftarrow \text{DPM-Solver}(\mathbf{x}^{(h)}, h, \mathbf{C}, \hat{S})$ 
8:      $\mathbf{x}^{(h-1)} \leftarrow m \odot \mathbf{x}_{\text{short}}^{(h-1)} + (1-m) \odot \mathbf{x}_{\text{branch}}^{(h-1)}$ 
9:   end for
10:   $B_{k,n} \leftarrow \mathbf{x}_{t_b:T}^{(0)}$ ,  $\Lambda_{k,n} \leftarrow [\hat{S}, B_{k,n}]$ 
11: end for
12: return  $\{\Lambda_{k,n}\}_{n=1}^N$ 

```

C. Pivot-Conditioned Diffusion Mechanism

We apply a binary temporal mask $m(t) = \mathbb{1}[t \leq t_b]$ during sampling, where $\mathbb{1}\{\cdot\}$ denotes the indicator; short-term frames ($t \leq t_b$) are frozen and long-term frames ($t > t_b$) are denoised. The operator $\text{re-noise}(\hat{S})$ re-applies the forward noising kernel to the clean short-term anchor \hat{S} at the current reverse step, i.e., it draws $\mathbf{x}_{\text{short}}^{(h-1)} \sim q_{h-1|0}(\cdot | \hat{S})$ and overwrites the short-term entries of the state.

Step A: Multiple Short-Term Segments Sampling and Pre-Filtering. In this part, we first sample K joint short-term segments $\{\hat{S}_k\} \sim p_\theta(S | \mathbf{C})$. For each candidate, we compute a PDM score over the short-term window $[0, t_b]$ that evaluates safety, rule compliance, comfort, and progress. If all scores are zero, the scene is deemed unsafe for nominal execution and an emergency-braking profile is triggered for this cycle. Otherwise, the remaining short-term segments are taken as anchors for the next steps.

Step B: Multi-Branch Completion Under Each Short-Term Segment. In this part, we complete diverse long-term branches conditioned on each retained shared short-term segment, with the goal of keeping identical near-term execution while exploring multiple plausible futures. For a retained short-term segment \hat{S}_k , we run masked reverse diffusion to generate the long-term branch while freezing all short-term frames, yielding N futures $\Lambda_{k,n} = [\hat{S}_k, B_{k,n}]$. The full sampling procedure is summarized in Algorithm 1, where the noise term ϵ_ϕ is derived from the model’s $\mathbf{x}^{(0)}$ estimate via the standard variance-preserving (VP) conversion; we set $\beta_h = 1 - \alpha_h$ and choose σ_h according to the VP schedule. At each reverse step, known short-term segments are re-noised from the anchor and overwrite the corresponding states; unknown branch frames are denoised by G_{full} and merged by the mask. This guarantees that all $\Lambda_{k,n}$ share the same near-term execution and branch only on $(t_b, T]$. As illustrated in Fig. 2, the known short-term segment is repeatedly re-noised to enforce consistency, and the unknown branch is generated by the model.

Step C: Ego Candidate Set Under a Shared Short-Term Segment. For each retained shared short-term segment \hat{S}_k , we build a separate ego candidate set $\mathcal{T}_{\text{short}}^{(k)}$. Within each set, all candidates are identical to the group’s anchor $\hat{\tau}_{0:t_b}^{0,k}$ on $[0, t_b]$ and branch only on $(t_b, T]$. Candidates are obtained exclusively

by reusing the ego branches produced by G_{full} under that short-term segment, yielding N alternative futures. The concrete scoring and selection procedure is described in Sec. IV-D.

D. Two-Stage Contingency-Aware Scoring

Stage 1: Short-Term Segment Pre-Filter ($[0, t_b]$). Within each short-term-segment group, all ego candidates share the same anchor on $[0, t_b]$; hence, the short-term score is a group-wise constant. For every retained short-term segment \hat{S}_k we compute

$$E_k^{\text{pre}} = \text{PDM}_{t_b}(\hat{\tau}_{0:t_b}^{0,k} | \hat{S}_k, C). \quad (9)$$

We retain only groups with $E_k^{\text{pre}} > 0$. If $\max_k E_k^{\text{pre}} = 0$, no short-term segment satisfies the executability criterion, and an emergency braking profile is issued for the current cycle.

Stage 2: Long-Term Window Multi-Scenario Aggregation.

For each surviving short-term segment k and any ego candidate $\tau^0 \in \mathcal{T}_{\text{short}}^{(k)}$, we evaluate the full-horizon PDM score on each generated future $\Lambda_{k,n}$:

$$E_{k,n}(\tau^0) = \text{PDM}_T(\tau^0 | \Lambda_{k,n}, C), \quad n = 1, 2, \dots, N, \quad (10)$$

where PDM_T denotes the PDM score over the full planning horizon $[0, T]$; in our experiments we instantiate it with $T = 8$ s (i.e., $\text{PDM}_{8\text{s}}$). We aggregate across scenarios by the mean

$$\bar{E}_k(\tau^0) = \frac{1}{N} \sum_{n=1}^N E_{k,n}(\tau^0). \quad (11)$$

The best candidate within the short-term segment k is then selected as

$$\tilde{\tau}^{0,k} = \arg \max_{\tau^0 \in \mathcal{T}_{\text{short}}^{(k)}} \bar{E}_k(\tau^0), \quad (12)$$

followed by cross-group selection across the short-term segments:

$$k^* = \arg \max_k \bar{E}_k(\tilde{\tau}^{0,k}), \quad \tau^{0,*} = \tilde{\tau}^{0,k^*}. \quad (13)$$

Stage 1 serves as an executability pre-filter, whereas Stage 2 selects the within-group and global optima by the mean PDM_T across generated futures.

V. EXPERIMENTS

A. Experimental Setup

Benchmark. We evaluate our method on the nuPlan benchmark [28], a large-scale dataset for closed-loop motion planning. It provides $\sim 1,500$ hours of expert driving across four cities and a simulator that initializes from logged scenarios and runs for 15 s at 10 Hz. We use the official simulator, which initializes scenarios from real-world logs. Other agents are controlled by either log-replay (non-reactive, **NR**) or an IDM policy (reactive, **R**). Low-level tracking is handled by an LQR controller that converts the planned trajectory into throttle/brake/steer commands. We report results on **Test14-hard** (272 scenarios) [29], **Test14** (261 scenarios), and **Val14** (1118 scenarios), each in both non-reactive and reactive modes.

Baselines and Metrics. We compare CoPlanner against the following representative planners: IDM [30]; PDM* (including PDM-Closed, PDM-Hybrid, and PDM-Open) [5]; UrbanDriver [31]; GameFormer [20]; PlanTF [29]; Diffusion Planner [32]; and PLUTO [33] (reported both with and without refine). Unless otherwise noted, we use official implementations and public settings. For the Diffusion Planner (learning-only) variant, we retrain on our unified $\sim 1\text{M}$ nuPlan split. For the variants with post-processing (refine), we report the publicly released results without modification. Following nuPlan, we report the overall **nuPlan score** (0–100, higher is better) averaged over scenarios, together with key sub-metrics: Collisions, Time-to-collision (TTC), Drivable area compliance, Comfort, and Progress.

B. Implementation Details

Data and Preprocessing. We train on $\sim 1\text{M}$ nuPlan clips with 2 s history and an 8 s planning horizon, sampled at 10 Hz. Poses are expressed in an ego-centric frame at $t=0$ with per-scene normalization. For each scene, we keep the ego vehicle and the 32 nearest other agents within a 100 m radius for context, while predicting only the 10 nearest agents. Agent histories include position and heading encoded as $(x, y, \cos \psi, \sin \psi)$. We also gather vectorized map and traffic-control elements and the route polyline within 100 m of the ego vehicle at the current time; neighboring agents' 2 s histories are stacked at 10 Hz. All modalities are padded to fixed token slots, and attention masks prevent attending to padded tokens or invalid segments. Following the time horizon used in PDM's hybrid setup and scoring scheme, we set the branching time to $t_b=4$ s.

Models. We use two DiT models operating on stacked multi-agent trajectory tokens with gated cross-attention to context tokens (vectorized lanes/traffic controls, route, and agent histories) and sinusoidal timestep embeddings. The short-term model G_{shared} generates a high-confidence shared short-horizon segment of length $t_b=4$ s; the full model G_{full} performs pivot-conditioned diffusion mechanism to produce diverse long-term branches conditioned on that short-term segment. Both models are trained with an $\mathbf{x}^{(0)}$ -prediction objective. For G_{full} , Gaussian noise is applied only to frames $t > t_b$ while short-term segment frames remain clean and are provided as conditions. To expose the model to full-trajectory multimodality, the x_0 target covers the entire 8 s horizon, and we apply short-term-conditioning dropout with a rate of 0.5 during training (randomly masking the short-term segment condition to reduce over-reliance). For sampling, we use DPM-Solver++ [34] with a VP noise schedule. We train for 300 epochs with a global batch size of 512 on $4 \times \text{RTX } 4090$ GPUs (about 2 days). The learning rate is initialized at 5×10^{-4} and uses a 5-epoch warmup.

CoPlanner. We implement CoPlanner on the nuPlan closed-loop benchmark. For candidate short-term segment generation, we retain $K = 3$ short-term joint trajectories $\{\hat{S}_k\}_{k=1}^K$, each validated by the PDM score. If all scores are zero, we trigger an emergency-braking profile. For each retained short-term segment \hat{S}_k , the pivot-conditioned diffusion mechanism samples $N = 3$ diverse long-term branches. Ego planning

TABLE I

CLOSED-LOOP PLANNING RESULTS ON nuPLAN DATASET. COLUMNS ARE ORDERED AS TEST14-HARD, TEST14, THEN VAL14. COLORS DENOTE THE TOP-3 WITHIN EACH GROUP (LEARNING-BASED VS. RULE-BASED & HYBRID): **FIRST**, **SECOND**, **THIRD**. **BOLD** = OVERALL BEST PER COLUMN (EXCLUDING EXPERT). NR: NON-REACTIVE; R: REACTIVE.

Type	Planner	Test14-hard		Test14		Val14	
		NR	R	NR	R	NR	R
Expert	Log-replay	85.96	68.80	94.03	75.86	93.53	80.32
Learning-based	PDM-Open	33.51	35.83	52.81	57.23	53.53	54.24
	UrbanDriver	50.40	49.95	51.83	67.15	68.57	64.11
	PlanTF	69.70	61.61	85.62	79.58	84.27	76.95
	PLUTO w/o refine	70.03	59.74	89.90	78.62	88.89	78.11
	Diffusion Planner w/o refine	72.39	61.21	87.33	75.51	86.87	75.30
	CoPlanner w/o refine (Ours)	76.82	64.47	90.31	78.81	89.48	79.00
Rule-based & Hybrid	IDM	56.15	62.26	70.39	74.42	75.60	77.33
	PDM-Closed	65.08	75.19	90.05	91.63	92.84	92.12
	PDM-Hybrid	65.99	76.07	90.10	91.28	92.77	92.11
	GameFormer	68.70	67.05	83.88	82.05	79.94	79.78
	PLUTO	80.08	76.88	92.23	90.29	92.88	76.88
	Diffusion Planner w/ refine	78.87	82.00	94.80	91.75	94.26	92.90
CoPlanner w/ refine (Ours)	76.05	78.59	93.47	92.00	94.45	93.13	

results are projected to a kinematic bicycle model to ensure trackability; a low-level LQR converts the selected trajectory into control commands.

VI. RESULTS AND DISCUSSION

All methods are evaluated in the official nuPlan closed-loop simulator in both non-reactive (log-replay) and reactive (IDM) settings with identical LQR tracking and controller limits to ensure fairness. To align with hybrid baselines, we attach a refinement module [5], [35] as the postprocessor to our approach (**CoPlanner w/ refine**).

A. Main Results

Table I summarizes closed-loop performance across **Test14-hard**, **Test14**, and **Val14** in both non-reactive (NR) and reactive (R) modes. On **Val14**, **CoPlanner w/ refine** achieves the best overall results in both NR and R; on **Test14**, **Diffusion Planner w/ refine** ranks first in NR while **CoPlanner w/ refine** ranks first in R and remains highly competitive in NR; on **Test14-hard**, **Diffusion Planner w/ refine** leads in R and is second in NR, with **CoPlanner w/ refine** close to the strongest hybrid baseline in NR. Notably, **CoPlanner w/o refine** already outperforms most learning-only baselines across splits, and the lightweight refinement consistently adds further gains. In this part, **CoPlanner w/o refine** and the single-stage Diffusion Planner share the same DiT backbone and training recipe; CoPlanner’s two-stage factorization lowers conditional entropy and simplifies credit assignment via shorter horizons, and reduces imitation bias. Overall, these results validate the benefit of CoPlanner’s shared short-term segment generation and multi-scenario scoring: it improves safety and comfort in interactive scenarios without sacrificing progress and remains robust on tail-risk scenarios. We observe comparatively smaller gains on **Test14-hard**, which we attribute to our deliberate reliance on a PDM-style composite score for pre-filtering and ranking: this split is curated around scenarios

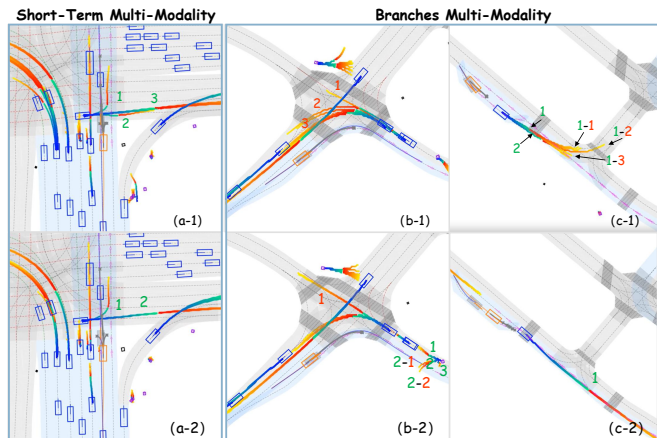


Fig. 3. Top: under a shared **short-term segment** (green), multiple **long-term branches** (red) emerge. Bottom: as time elapses, branches are pruned and converge to one or a few dominant outcomes.

where PDM scoring is known to be less informative, reducing the advantage of our scorer.

B. Qualitative Results

Figure 3 illustrates CoPlanner’s short-term conditioned behavior across three interaction scenes. In the top row, trajectories share a common green short-term segment and then branch into multiple red branches, capturing early-stage multimodality at merges, lane changes, and unprotected turns. In the bottom row, as time elapses and interactions resolve, many less likely branches are pruned by right of way, map topology, and social constraints, and the set converges to one or a few dominant outcomes. The shared short-term segment yields stable near-term execution and prevents plan switching, while the surviving branches keep feasible contingency options if surrounding agents brake, force their way in, or yield unpredictably. This leads to smoother and safer motion without sacrificing route progress.

TABLE II
ABLATIONS ON NUPLAN VAL14 (REACTIVE). HIGHER IS BETTER.

Block	Variant	Eval. Hori.	Score	Collisions	TTC	Drivable	Comfort	Progress
Planner Arch. w/o refine	Diffusion Planner	-	75.30	95.57	91.50	96.69	85.50	89.17
	CoPlanner	-	79.00	96.14	92.27	97.66	88.14	90.84
Evaluation Horizon	PDM Score	8s	91.97	98.73	95.20	99.64	89.04	98.37
	PDM Score	4s	92.95	98.21	93.38	99.73	92.67	99.64
Generation Strategy (Contingency Score)	Regular Sampling	8s	92.89	98.26	93.38	99.73	91.68	99.64
	Direct Assignment	8s	93.03	98.35	94.63	99.82	94.19	99.02
	Ours	8s	93.13	98.26	94.81	99.73	94.28	99.19

C. Ablation Studies

We structure the ablations into three parts, summarized in Table II and illustrated in Fig. 4: (i) architecture compares single-stage diffusion with our two-stage generation without refinement; (ii) scoring baseline keeps the refinement module and extends the PDM evaluation window to 8 s; (iii) our method: Contingency Score with the same refinement and 8 s window, comparing three sampling strategies.

Planner Architecture. In this experiment, all models are trained on the same dataset. The two-stage design with a shared short-term segment followed by branch completion improves the nuPlan score by 3.7 points over the single-stage baseline and raises every submetric. This suggests that anchoring a validated shared short-term segment stabilizes near-term execution and yields more socially consistent interactions. Methodologically, decoupling the problem into a short-horizon executor and a short-term segment conditioned long-term completer reduces the effective learning horizon faced by imitation learning, mitigates imitation shift, and results in stronger closed-loop behavior.

Evaluation Horizon. With the same refinement module, extending the PDM evaluation window from 4 s to 8 s lowers the overall nuPlan score. Safety metrics improve (fewer collisions, higher TTC), but comfort degrades noticeably and progress shows no gain. The longer window overemphasizes distant, hypothetical interactions and underweights near-term executability, prompting overly conservative behavior (earlier braking and jerkier control). This observation motivates our design: score the shared short-term segment on $[0, t_b]$ and aggregate long-term costs across multiple futures rather than stretching a single-window PDM.

Generation Strategy Under Contingency Score. With the same refinement and 8 s window, Contingency Score combined with pivot-conditioned diffusion mechanism performs best among the three strategies considered. Internally, we compare three sampling schemes in Fig. 4: (a) Regular sampling increases diversity but suffers from uncertainty accumulation over the long 8 s horizon, which scatters modes and weakens temporal coherence for immediate decision-making; (b) Frame-wise short-term segment direct assignment reduces dispersion but distorts the input distribution through frequent re-anchoring, leading to jitter and mode flips; (c) Pivot-conditioned diffusion mechanism maintains a consistent shared short-term segment before branching, which preserves the underlying data distribution while balancing stability and

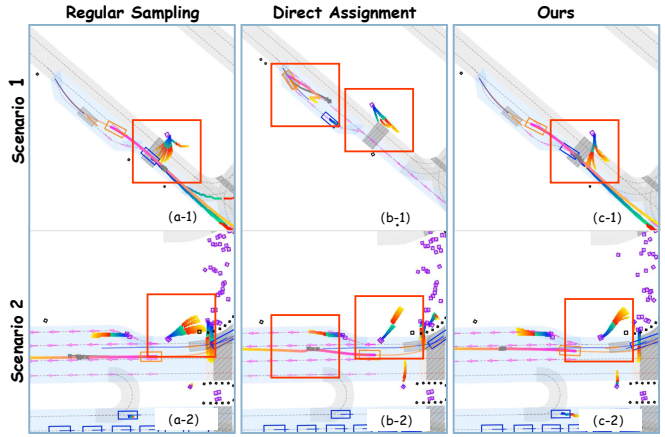


Fig. 4. Comparison of three generation strategies for multi-modal trajectory generation. (a) Regular Sampling: multiple futures are generated independently per step, leading to scattered modes and temporal inconsistency. (b) Frame-wise short-term segment direct assignment: re-assigns the short-term segment at every frame, which reduces spread but causes jitter and mode flipping. (c) Ours: short-term segment consistent multi-modal prediction with late-stage branching, producing a stable shared short-term segment and diverse yet feasible terminations with improved temporal consistency.

diversity, achieving the best overall performance. In particular, pivot-conditioned diffusion mechanism achieves the highest scores in both safety and comfort metrics, leading to the best overall performance. Further, compared across blocks under the same 8 s setting, pivot-conditioned diffusion mechanism with Contingency Score surpasses the PDM baseline and also slightly edges the 4 s PDM setup overall, chiefly by improving comfort and TTC. This indicates that our method better balances near-term executability and long-term preparedness.

VII. CONCLUSION AND FUTURE WORK

In this work, we present *CoPlanner*, a unified generation–evaluation framework that couples short-term conditioned diffusion with contingency-aware scoring. By enforcing a safety-validated shared short-term segment and sampling diverse long-term branches, *CoPlanner* stabilizes near-term execution while retaining feasible fallback options when interactions evolve. On the nuPlan closed-loop benchmark, *CoPlanner* delivers consistent gains: it attains state-of-the-art performance on **Val14** in both non-reactive and reactive modes and on **Test14** (reactive), and remains competitive on the other splits. Ablations indicate that (i) the two-stage generation with a shared short-term segment and (ii) the multi-scenario aggregation in the evaluator are the primary drivers

of the improvement, yielding better safety and comfort without sacrificing progress.

Limitations and Future Work. CoPlanner may still fail when all sampled futures are infeasible, and the current aggregator does not model scenario likelihoods or calibrated uncertainty. The branching time t_b is fixed, which can be suboptimal across scenes, and the computational budget for multi-branch evaluation can be further reduced. Future work will add feasibility-aware guidance during denoising, incorporate physics-informed and dynamics-informed priors, and adaptively choose the branching horizon. It will also aggregate scenarios using likelihood or risk weights under calibrated uncertainty, and broaden closed-loop evaluations with stronger reactive agents and real-world tests.

REFERENCES

- [1] S. Hagedorn, M. Hallgarten, M. Stoll, and A. P. Condurache, “The Integration of Prediction and Planning in Deep Learning Automated Driving Systems: A Review,” *IEEE Transactions on Intelligent Vehicles*, vol. 10, no. 5, pp. 3626–3643, 2025.
- [2] R. Yao and X. Sun, “Hierarchical Prediction Uncertainty-Aware Motion Planning for Autonomous Driving in Lane-Changing Scenarios,” *Transportation Research Part C: Emerging Technologies*, vol. 171, p. 104962, 2025.
- [3] Z. Huang, Z. Zhang, A. Vaidya, Y. Chen, C. Lv, and J. F. Fisac, “Versatile Behavior Diffusion for Generalized Traffic Agent Simulation,” *arXiv preprint arXiv:2404.02524*, 2024.
- [4] C. Yang, Y. He, A. X. Tian, D. Chen, J. Wang, T. Shi, A. Heydariyan, and P. Liu, “WcDT: World-Centric Diffusion Transformer for Traffic Scene Generation,” in *International Conference on Robotics and Automation*, pp. 6566–6572, IEEE, 2025.
- [5] D. Dauner, M. Hallgarten, A. Geiger, and K. Chitta, “Parting with Misconceptions about Learning-based Vehicle Motion Planning,” in *Conference on Robot Learning*, 2023.
- [6] Z. Huang, X. Weng, M. Igl, Y. Chen, Y. Cao, B. Ivanovic, M. Pavone, and C. Lv, “Gen-Drive: Enhancing Diffusion Generative Driving Policies with Reward Modeling and Reinforcement Learning Fine-tuning,” *arXiv preprint arXiv:2410.05582*, 2024.
- [7] A. Cui, S. Casas, A. Sadat, R. Liao, and R. Urtasun, “LookOut: Diverse Multi-Future Prediction and Planning for Self-Driving,” in *International Conference on Computer Vision*, pp. 16107–16116, 2021.
- [8] K. A. Mustafa, D. J. Ornia, J. Kober, and J. Alonso-Mora, “RACP: Risk-Aware Contingency Planning With Multi-Modal Predictions,” *IEEE Transactions on Intelligent Vehicles*, 2024.
- [9] M.-K. Bouzidi, B. Derajic, D. Goehring, and J. Reichardt, “Reachability-Based Contingency Planning against Multi-Modal Predictions with Branch MPC,” *arXiv preprint arXiv:2502.02550*, 2025.
- [10] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, “Trajectron++: Dynamically-Feasible Trajectory Forecasting with Heterogeneous Data,” in *European Conference on Computer Vision*, pp. 683–700, Springer, 2020.
- [11] Y. Chen, J. Cheng, L. Gan, S. Wang, H. Liu, X. Mei, and M. Liu, “IR-STP: Enhancing Autonomous Driving With Interaction Reasoning in Spatio-Temporal Planning,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 8, pp. 10331–10343, 2024.
- [12] H. Liu, L. Chen, Y. Qiao, C. Lv, and H. Li, “Reasoning Multi-Agent Behavioral Topology for Interactive Autonomous Driving,” *arXiv preprint arXiv:2409.18031*, 2024.
- [13] J. Gao, C. Sun, H. Zhao, Y. Shen, D. Anguelov, C. Li, and C. Schmid, “VectorNet: Encoding HD Maps and Agent Dynamics From Vectorized Representation,” in *IEEE Conference on Computer Vision and Pattern Recognition*, June 2020.
- [14] S. Shi, L. Jiang, D. Dai, and B. Schiele, “MTR++: Multi-Agent Motion Prediction With Symmetric Scene Modeling and Guided Intention Querying,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 46, no. 5, pp. 3955–3971, 2024.
- [15] T. Phan-Minh, E. C. Grigore, F. A. Boulton, O. Beijbom, and E. M. Wolff, “CoverNet: Multimodal Behavior Prediction Using Trajectory Sets,” in *IEEE Conference on Computer Vision and Pattern Recognition*, June 2020.
- [16] N. Lee, W. Choi, P. Vernaza, C. B. Choy, P. H. S. Torr, and M. Chandraker, “DESIRE: Distant Future Prediction in Dynamic Scenes With Interacting Agents,” in *IEEE Conference on Computer Vision and Pattern Recognition*, July 2017.
- [17] Q. Sun, X. Huang, J. Gu, B. C. Williams, and H. Zhao, “M2I: From Factored Marginal Trajectory Prediction to Interactive Prediction,” in *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6543–6552, 2022.
- [18] B. Ma, W. Liu, and J. Ma, “Trajectory Tree-Based Pairwise Game for Interactive Decision-Making and Motion Planning in Autonomous Driving,” *IEEE Transactions on Vehicular Technology*, pp. 1–13, 2025.
- [19] Z. Huang, Y. Xie, B. Ma, S. Shen, and J. Ma, “Fast and Scalable Game-Theoretic Trajectory Planning with Intentional Uncertainties,” *arXiv preprint arXiv:2507.12174*, 2025.
- [20] Z. Huang, H. Liu, and C. Lv, “GameFormer: Game-theoretic Modeling and Learning of Transformer-based Interactive Prediction and Planning for Autonomous Driving,” in *International Conference on Computer Vision*, pp. 3903–3913, 2023.
- [21] H. Liu, Z. Huang, W. Huang, H. Yang, X. Mo, and C. Lv, “Hybrid-Prediction Integrated Planning for Autonomous Driving,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 47, no. 4, pp. 2597–2614, 2025.
- [22] R. Yao, Y. Wang, H. Liu, R. Yang, Z. Peng, L. Zhu, and J. Ma, “CALMM-Drive: Confidence-Aware Autonomous Driving with Large Multimodal Model,” *arXiv preprint arXiv:2412.04209*, 2024.
- [23] X. Chen, J. Yan, W. Liao, T. He, and P. Peng, “Int2Planner: An Intention-based Multi-modal Motion Planner for Integrated Prediction and Planning,” *arXiv preprint arXiv:2501.12799*, 2025.
- [24] M. Zheng, L. Zheng, L. Zhu, and J. Ma, “Occlusion-Aware Consistent Model Predictive Control for Robot Navigation in Occluded Obstacle-Dense Environments,” *arXiv preprint arXiv:2503.04563*, 2025.
- [25] L. Zheng, R. Yang, M. Zheng, M. Y. Wang, and J. Ma, “Safe and Real-Time Consistent Planning for Autonomous Vehicles in Partially Observed Environments via Parallel Consensus Optimization,” *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–17, 2025.
- [26] Y. Chen, P. Karkus, B. Ivanovic, X. Weng, and M. Pavone, “Tree-structured Policy Planning with Learned Behavior Models,” in *IEEE International Conference on Robotics and Automation*, pp. 7902–7908, 2023.
- [27] 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 *IEEE International Conference on Robotics and Automation*, pp. 6806–6812, 2024.
- [28] N. Karnchanachari, D. Geromichalos, K. S. Tan, N. Li, C. Eriksen, S. Yaghoubi, N. Mehdipour, G. Bernasconi, W. K. Fong, Y. Guo, et al., “Towards learning-based planning: The nuPlan benchmark for real-world autonomous driving,” in *International Conference on Robotics and Automation*, pp. 629–636, 2024.
- [29] J. Cheng, Y. Chen, X. Mei, B. Yang, B. Li, and M. Liu, “Rethinking Imitation-based Planners for Autonomous Driving,” in *International Conference on Robotics and Automation*, pp. 14123–14130, 2024.
- [30] D. Helbing and B. Tilch, “Generalized force model of traffic dynamics,” *Physical review E*, vol. 58, no. 1, p. 133, 1998.
- [31] O. Scheel, L. Bergamini, M. Wolczyk, B. Osiniński, and P. Ondruska, “Urban driver: Learning to drive from real-world demonstrations using policy gradients,” in *Conference on Robot Learning*, pp. 718–728, PMLR, 2022.
- [32] Y. Zheng, R. Liang, K. Zheng, J. Zheng, L. Mao, J. Li, W. Gu, R. Ai, S. E. Li, X. Zhan, and J. Liu, “Diffusion-Based Planning for Autonomous Driving with Flexible Guidance,” in *International Conference on Learning Representations*, 2025.
- [33] 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.
- [34] C. Lu, Y. Zhou, F. Bao, J. Chen, C. Li, and J. Zhu, “DPM-Solver: A Fast ODE Solver for Diffusion Probabilistic Model Sampling in Around 10 Steps,” *Advances in neural information processing systems*, vol. 35, pp. 5775–5787, 2022.
- [35] Q. Sun, H. Wang, J. Zhan, F. Nie, X. Wen, L. Xu, K. Zhan, P. Jia, X. Lang, and H. Zhao, “Generalizing Motion Planners with Mixture of Experts for Autonomous Driving,” in *International Conference on Robotics and Automation*, pp. 6033–6039, 2025.