

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

Hi-Drive: Hierarchical POMDP Planning for Safe Autonomous Driving in Diverse Urban Environments

Xuanjin Jin, Chendong Zeng, Shengfa Zhu, Chunxiao Liu, Panpan Cai

Abstract—Uncertainties in dynamic road environments pose significant challenges for behavior and trajectory planning in autonomous driving. This paper introduces Hi-Drive, a hierarchical planning algorithm addressing uncertainties at both behavior and trajectory levels using a hierarchical Partially Observable Markov Decision Process (POMDP) formulation. Hi-Drive employs driver models to represent uncertain behavioral intentions of other vehicles and uses their parameters to infer hidden driving styles. By treating driver models as high-level decision-making actions, our approach effectively manages the exponential complexity inherent in POMDPs. To further enhance safety and robustness, Hi-Drive integrates a trajectory optimization based on importance sampling, refining trajectories using a comprehensive analysis of critical agents. Evaluations on real-world urban driving datasets demonstrate that Hi-Drive significantly outperforms state-of-the-art planning-based and learning-based methods across diverse urban driving situations in real-world benchmarks.

Index Terms—Autonomous driving, Uncertainty handling, Partially Observable Markov Decision Process (POMDP), Hierarchical Planning, Behavior and Trajectory Planning

I. INTRODUCTION

Autonomous driving (AD) aims to improve transportation convenience by enabling vehicles to sense, plan and act autonomously. A critical component of AD is behavior and trajectory planning, which determines high-level driving behaviors (lane-following, lane-changing, turning, etc.) and low-level trajectories (specifying positions, speeds, and accelerations) of the autonomous vehicle (*ego-vehicle*). These decisions must be carefully made to ensure safety and efficiency.

A significant challenge in behavior and trajectory planning arises from uncertainty, particularly regarding unknown behavioral intentions and driving styles of other traffic participants (*exo-agents*). For example, subtle actions of nearby vehicles, such as slowly drifting towards the lane boundary, can reflect diverse intentions. These could range from preparing to change lanes or trying to create space for another vehicle, to simply exhibiting reckless driving habits. Hedging against such uncertainties is essential for safe and robust planning.

Manuscript received: April 8, 2025; Revised: August 7, 2025; Accepted: September 24, 2025. This paper was recommended for publication by Editor Ashis Banerjee upon evaluation of the Associate Editor and Reviewers' comments. This work was supported in part by the National Natural Science Foundation of China under Grant 62303304. (*Corresponding Author: Panpan Cai.*)

Xuanjin Jin, Chendong Zeng and Panpan Cai are with Shanghai Jiao Tong University, Shanghai 201100, China, and also with Shanghai Innovation Institute, Shanghai 200030, China (e-mail: xuanjin.jin@sjtu.edu.cn; chendongzeng@sjtu.edu.cn; cai_panpan@sjtu.edu.cn).

Shengfa Zhu and Chunxiao Liu are with SenseTime, Shanghai 200240, China. (e-mail: zhushengfa@senseauto.com; liuchunxiao@senseauto.com).

Digital Object Identifier (DOI): see top of this page.

©2026 IEEE

Optimally addressing these uncertainties is computationally intractable in real-time settings, due to the exponential growth of computation with the number of traffic participants and the planning horizon, known respectively as the “curse of dimensionality” and the “curse of history” [1]. To mitigate complexity, existing approaches typically simplify uncertainty modeling by addressing only specific aspects, such as behavior-level uncertainty [2]–[12], trajectory-level uncertainty [13]–[16], or by employing maximum-likelihood assumptions [17]–[20]. However, these simplifications can compromise safety and robustness, especially in scenarios involving complex interactions among vehicles.

In this work, we propose Hi-Drive, a hierarchical POMDP planner that, for the first time, enables joint reasoning over uncertainties in both high-level behavioral intentions and low-level driving styles for general urban driving. We treat the behavioral intentions and driving styles of *exo-agents* as latent states, represented using driver models and their free parameters, respectively. We develop a hierarchical Bayesian filter to track a hierarchy of probability distributions over these hidden states (*hierarchical belief*) over time, enabling structured and efficient reasoning about complex and uncertain interactions.

To address the computational complexity inherent to POMDP planning, Hi-Drive also uses high-level driver models to represent actions for the *ego-vehicle*, thereby shortening the effective planning horizon. At the lower level, driving trajectories are generated by simulating these driver models, allowing joint evaluation of decisions across both behavior and trajectory levels. Additionally, to enhance trajectory robustness, a subsequent trajectory optimization step employs importance sampling, re-generating and re-evaluating trajectories, with a focus on critical agents.

Experimental evaluation shows the effectiveness of Hi-Drive for safe and robust long-term planning in urban environments. Evaluations on the large-scale nuPlan benchmark [21] show that Hi-Drive achieves competitive performance compared to learning-based approaches, despite requiring no training. Further analysis on an additional real-world dataset, featuring long-term tasks, shows substantial improvements in safety and task completion compared to existing planning-based methods. We attribute Hi-Drive’s advantage to hierarchical and long-term reasoning and the application of importance sampling. Furthermore, Hi-Drive can perform real-time planning in complex driving scenes, even with limited computational resources. It also allows transparent interpretation of its reasoning, as shown in the supplementary video.

In summary, our main contributions include:

- A hierarchical POMDP planner that synergistically and efficiently handles uncertainties at both behavior and

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

trajectory levels.

- Integration of POMDP planning with an importance-sampling-based trajectory optimizer to enhance trajectory safety and robustness.
- This work is the first to apply POMDP planning to general urban scenes (instead of pre-selected scenes such as multi-lane roads) and achieve state-of-the-art performance in large-scale benchmarks.

II. RELATED WORK

Safe and robust autonomous driving requires hedging against uncertainties at both behavior and trajectory levels. High-level uncertainties involve predicting others' routes, maneuvers, or interaction decisions, while low-level uncertainties correspond to variations in acceleration, steering, etc. However, optimally tackling both levels of uncertainty leads to intractable computational complexity, thus, it requires simplification to achieve real-time performance. Planning under uncertainty is often simplified to make it computationally tractable. A straightforward approach is to plan for the maximum-likelihood behaviors of others, commonly using Monte-Carlo Tree Search (MCTS) [17]–[20]. While efficient, this may pose risks in scenes with low-probability events.

Other methods address the behavior-level uncertainty or trajectory-level uncertainty in isolation. This separation, coupled with prior focus on specific environments like highways, compromises safety in complex and general urban scenarios where both intentions and styles interact. For instance, to handle trajectory uncertainty, planners use algorithms like RRT* or numerical optimization, but these are often decoupled from a high-level behavioral planner, which prohibits synergistic behaviors like changing lanes to avoid an unpredictable driver [13]–[16]. To handle behavior uncertainty, Multipolicy Decision-Making (MPDM) evaluates a set of candidate policies for the ego-vehicle [2]–[4]. However, its open-loop nature ignores the benefits of future information, which can lead to overly conservative actions.

To perform closed-loop planning under uncertainty, past works formulate the driving problem as POMDPs, treating the unknown behaviors of others as hidden states, then apply POMDP solvers such as QMDP [12], POMCP [6]–[8], DESPOT [10], [11], and ABT [9] to compute the ego-vehicle's behavior policy conditioned on future observations. Some of these works perform short-term planning by searching over primitive actions, such as longitudinal accelerations and lateral velocities [5], [6], [8]–[11], [22], while others consider high-level behaviors such as lane keeping and changing for long-term planning [7], [12]. However, existing POMDP methods are typically limited to specific scenes, such as multi-lane roads. Moreover, they only address behavior-level uncertainty and neglect trajectory-level uncertainty for computational simplicity. This is insufficient given that different drivers with the same intention may exhibit vastly different trajectories—e.g., an aggressive driver might cut into a lane abruptly, while a conservative driver would do so more cautiously. In this work, we extend the POMDP approach to tackle urban environments and achieve state-of-the-art results on a large-scale benchmark,

nuPlan [21]. We also use a hierarchical POMDP framework to address uncertainty at both behavior and trajectory levels simultaneously, and complement it with a trajectory optimizer based on importance sampling to further enhance driving safety and robustness within limited planning time.

Our work also advances on prior hierarchical and factored POMDPs by introducing a synergistic framework that, unlike sequential switching of granularity [23], [24], jointly reasons about high-level intentions and low-level styles. We also use an agent-centric factorization of hidden states that, in contrast to state-region, observability, or objective-based decompositions [25]–[28], ensures scalability with the number of agents.

III. OVERVIEW

In this paper, we propose Hi-Drive, a hierarchical planner (Fig. 1) addressing uncertainties at both behavior and trajectory levels. Our approach is grounded in a Partially Observable Markov Decision Process (POMDP) framework, modeling behavioral intentions and driving styles of exo-agents as hidden states, and treating the behavioral intention of the ego-vehicle as actions.

Hi-Drive consists of three interconnected modules:

Multi-model Inference (Fig. 1a): Behavioral intentions of exo-agents are represented by distinct driver models, while driving styles are characterized by their parameters. We employ hierarchical Bayesian inference—combining exact filtering at the intention level and particle filtering at the style and state levels—to dynamically update and track beliefs about exo-agents' behaviors based on observations.

Belief Tree Search (Fig. 1b): Conditioned on the updated beliefs, we perform an online belief tree search to identify the optimal high-level behavior for the ego-vehicle. The algorithm samples *driving scenarios* (see Section VI-A) from the current belief, then recursively explores possible actions and resulting observations to estimate their future outcomes and uncertainties. An optimal behavior sequence is thus determined from this belief-space planning process.

Trajectory Optimization (Fig. 1c): Following the selected optimal behavior sequence, we apply importance sampling to resample representative driving scenarios, emphasizing the analysis of critical agents. These sampled scenarios are then used to refine low-level driving trajectories for the ego-vehicle. Finally, a cross-scenario evaluation selects the most robust and effective trajectory for execution.

IV. HIERARCHICAL POMDP MODELING

Our approach is centered on a hierarchical POMDP formulation, using high-level driver models as its building blocks.

A. Driver Models

We consider two primary driving behaviors: lane-following (*LF*) and lane-changing (*LC*). Each behavior is represented by a driver model:

- *Lane-following (LF)* guides vehicles along the centerlines of lanes and through lane connectors at intersections.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

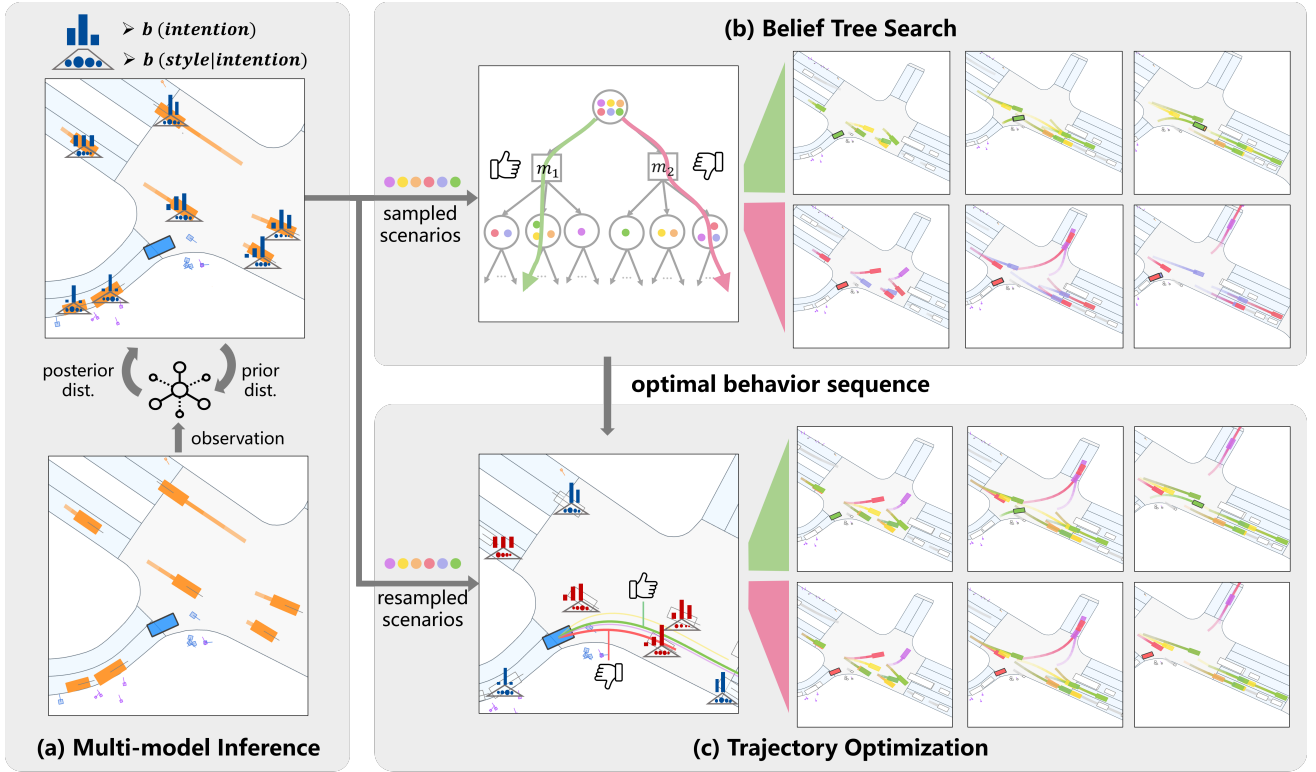


Fig. 1: Overview of Hi-Drive: (a) The multi-model inference tracks beliefs about exo-agents’ behavioral intentions and driving styles based on real-time observations. (b) The belief tree search determines the ego-vehicle’s optimal policy over high-level behaviors. (c) Trajectory optimization resamples representative scenarios using importance sampling, cross-evaluates low-level trajectories generated by those scenarios, and selects the one with the best driving performance. See detailed explanations of the driving scene visualization in Section VIII-E.

- *Lane-changing (LC)* uses the MOBIL decision-making model [29], which ensures sufficient gap and minimal interference with the target lane.

Both behaviors employ pure pursuit [30] for lateral trajectory control and the Intelligent Driver Model (IDM) [31] for longitudinal acceleration control. Driving style differs by behavior: for lane-following, style is represented by IDM’s desired speed (higher values indicate greater aggressiveness); for lane-changing, it is characterized by the look-ahead distance in pure pursuit (larger distances indicate more cautious driving).

Multiple *LF* or *LC* options can exist simultaneously, e.g., vehicles can select different outgoing lanes at intersections. To reduce computational complexity, we dynamically filter behaviors to include only legal actions consistent with the road network topology.

B. POMDP Model

The POMDP model is formulated as follows:

1) *State and Observation*: A state in our POMDP formulation includes:

- Behavioral intentions (driver models) of exo-agents: $m = \{m_i\}_{i \in I_{\text{exo}}}$.
- Driving styles (model parameters) of exo-agents: $\theta = \{\theta_i\}_{i \in I_{\text{exo}}}$.
- Physical states of all agents: $x = \{x_{\text{ego}}\} \cup \{x_i\}_{i \in I_{\text{exo}}}$, comprising positions, speeds, headings, etc.

Observations $o = \{o_{\text{ego}}\} \cup \{o_i\}_{i \in I_{\text{exo}}}$ reflect noisy measurements of the agents’ physical states.

2) *Action Space*: Actions for the ego-vehicle are high-level behaviors defined also by the driver models (*LFs*, *LCs*). Formally, an action is $a = (m_{\text{ego}}, \theta_{\text{ego}})$, where m_{ego} represents the behavior to decide for the ego-vehicle, and θ_{ego} is a pre-defined style value.

3) *State Transition and Observation Models*: The state transition function is deterministic and given by $x^t = G(x^{t-1}, m, \theta, m_{\text{ego}}, \theta_{\text{ego}})$, where G simulates one step forward using the specified driver models and style parameters. The observation model assumes Gaussian noise around the true physical states: $p(o^t | x^t) = \mathcal{N}(o^t; x^t, \Sigma)$.

4) *Reward Function*: The reward function evaluates key performance aspects:

- *Safety*: Cubic penalty for collisions, scaled by speed.
- *Efficiency*: Linear penalty based on deviation from desired speed.
- *Task Completion*: Exponential penalty for deviation from the target lane.
- *Smoothness*: Constant penalty for each lane change.

V. BELIEF TRACKING WITH MULTI-MODEL INFERENCE

This section presents multi-model inference for effectively tracking exo-agents’ behavioral intentions and driving styles within the POMDP framework.

A. Independence Assumptions

Direct Bayesian inference over the high-dimensional hidden states of our POMDP, which include N discrete intention variables and approximately $N \times M$ continuous style parameters (for N exo-agents and M intentions per agent), is computationally prohibitive. To achieve real-time performance, we employ two conditional independence assumptions:

- Behavioral intentions of exo-agents are assumed independent during the inference window, allowing independent inference of the N intention variables.
- Driving styles are assumed independent across different exo-agents and intentions, enabling independent inference of the $N \times M$ style parameters.

These assumptions yield a factored belief structure that significantly reduces inference complexity, making it linear with respect to the number of agents and intentions. Note that the independence applies *only* to latent intentions and style parameters. Contextual interactions between agents remain fully modeled through the resulting behaviors, e.g., a *LC* behavior falls back to *LF* when the gap in the target lane is insufficient.

Besides, we assume intentions and styles are static within each planning cycle. Although dynamic intentions could be modeled using a learned transition function, our experiments suggest that, in practice, it has minimal benefit under computational constraints.

B. Hierarchical Bayesian Filter

We maintain the factorized belief using hierarchical Bayesian filtering, consisting of low-level particle filters and high-level exact filters.

1) *Low-Level Particle Filter*: For each intention hypothesis $m_{i,k}$ of exo-agent i , we run a particle filter [32] to approximate the belief over driving style θ_i and physical state x_i^t :

$$b^t(x_i^t, \theta_i | m_i = m_{i,k})$$

Each particle filter consists of a *prediction step*, propagating particles forward via deterministic transition, $x_i^t = G_i(x_i^{t-1}, m_i, \theta_i, m_{\text{ego}}, \theta_{\text{ego}})$, and an *update step* adjusting particle weights based on observation likelihoods, $p(o_i^t | x_i^t) = \mathcal{N}(o_i^t; x_i^t, \Sigma)$.

2) *High-Level Exact Filter*: We apply exact Bayesian filtering at the intention level. The belief update for exo-agent i 's intention m_i is computed as:

$$b^t(m_i) = \eta p(o_i^t | m_i) b^{t-1}(m_i),$$

where η is a normalization constant.

The challenge is to efficiently calculate the observation likelihood $p(o_i^t | m_i)$. We derive:

$$p(o_i^t | m_i) = \sum_{x_i^{t-1}, \theta_i} p(o_i^t | G_i(x_i^{t-1}, m_i, \theta_i)) p(x_i^{t-1}, \theta_i | m_i).$$

This equation suggests upward information propagation, integrating observation likelihoods of low-level particles at the current step, $p(o_i^t | G_i(x_i^{t-1}, m_i, \theta_i))$, with particle weights from the last step, $p(x_i^{t-1}, \theta_i | m_i)$, to compute the high-level likelihood. Reuse of information enables efficient inference.

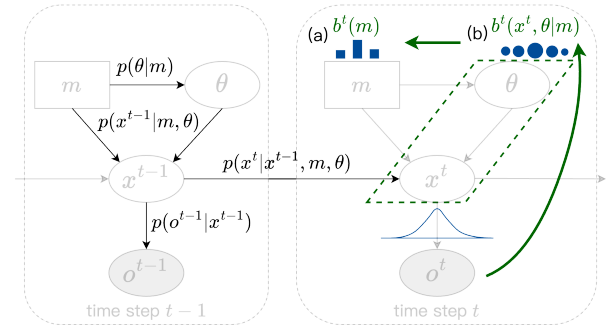


Fig. 2: The hierarchical belief tracker. Green arrows represent the process of updating the belief. (a) The belief over behavioral intention m . (b) The belief over physical states x^t and driving styles θ .

VI. SCENARIO-BASED BELIEF TREE SEARCH

Conditioned on the current belief, our belief tree search performs hierarchical behavior and trajectory planning, based on the POMDP formulation. The search algorithm is built upon DESPOT [33], an online POMDP planner that guarantees asymptotic optimality, extended here to support hierarchical beliefs and varying-length actions.

A. Sampling Driving Scenarios

The root belief node contains sampled *driving scenarios*. Each scenario ϕ is generated by sampling behavioral intentions (driver models, m) from the high-level exact filter and driving styles (model parameters, θ) and physical states (x) from low-level particle filters. These are combined with a sequence of random numbers $(\psi_t, \psi_{t+1}, \dots, \psi_{t+D})$ to determinize future actions and observations, where D is the maximum search depth. Formally, a scenario is expressed as $\phi = (x, m, \theta, \psi_t, \psi_{t+1}, \dots, \psi_{t+D})$, succinctly encoding hypothetical dynamics of the driving scene using N intention values, N style values and D random numbers. The belief tree is constructed based on K scenarios to ensure fast computation, while maintaining asymptotic optimality.

B. Abstract Action and Observation Branching

Action branching occurs when expanding each belief node, representing the ego-vehicle's selectable driving behaviors. The action space is dynamically constructed from the union of legal behaviors available to scenarios within that belief. Each action—such as a 2-second Lane-Follow (LF) or a 4-second Lane-Change (LC)—advances its corresponding scenarios using the forward simulation model $x^t = G(x^{t-1}, m, \theta, m_{\text{ego}}, \theta_{\text{ego}})$. Each completed action creates a new action node containing the updated scenarios. It also carries a step reward calculated as the cumulative reward averaged over the encountered scenarios.

Observation branching occurs under each action node, reflecting that ego-vehicle's observations may differ under various scenarios. Observations are deterministically generated using the *determinized* observation function $Z(o_t | x_t, \psi_t)$, by sampling Gaussian noise in $p(o_t | x_t)$ with ψ_t . Continuous

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

observations are uniformly discretized, forming abstract observation branches. Each branch creates new a belief node containing a subset of scenarios inherited from the parent.

Abstract branching greatly reduced the search complexity, while the hierarchical belief structure maintained in belief nodes enables coupled evolution of behaviors and trajectories for all agents.

C. Iterative Tree Construction

The belief tree iteratively expands via exploration paths. The exploration adopt DESPOT’s heuristics, selecting the action branch with the highest upper bound and the observation branch with the highest Weighted Excess Uncertainty (WEU, see [33]). It expands encountered leaf nodes, and proceeds until a maximum look-ahead horizon (e.g., 9 seconds) is reached, or the WEU at the leaf node becomes negative. Heuristic values are computed for the terminal leaf node. Afterwards, the search starts a *value backup* process towards the root, updating the values of all nodes throughout the path, following the Bellman’s principle. Upon reaching the root, a new exploration path starts.

D. Extracting Optimal Behavior Sequence

The belief tree search terminates when a planning time limit (e.g., 0.1 second) is reached. Unlike typical online planners that output only the immediate optimal action, we extract an optimal action sequence from the tree, to facilitate long-term trajectory planning. This is done by traversing the highest-value action branches and the most-likely observation branches.

This procedure also generates driving trajectories for scenarios encountered under the path. However, these trajectories have not been reliably evaluated due to sparsity of the belief tree. It thus requires an additional trajectory optimization module for robust planning.

VII. CROSS-SCENARIO TRAJECTORY OPTIMIZATION WITH IMPORTANCE SAMPLING

The cross-scenario trajectory optimization conducts a more detailed analysis at the trajectory level, thereby enhancing robustness to uncertainty.

Scenario Resampling: We employ importance sampling (IS) to resample driving scenarios, focusing specifically on *critical agents*, defined as those near the ego-vehicle or whose predicted paths intersect with the planned trajectory. Intentions of critical agents are uniformly resampled to ensure extensive coverage, while intentions and styles of non-critical agents are sampled directly from the current belief.

Trajectory Regeneration: For each resampled scenario, we regenerate the driving trajectory of the ego-vehicle by forward simulating the optimal behavior sequence determined by the belief tree search. Exo-agents are simulated simultaneously based on their sampled intentions and driving styles.

Cross Evaluation: Each candidate trajectory for the ego-vehicle is cross-evaluated against all resampled scenarios. The one most robust under uncertainty is selected. Specifically,

for each cross-evaluation trajectory-scenario pair, we simulate forward the world state conditioned on the ego-vehicle’s candidate trajectory. The resulting state sequence is evaluated using the reward model (Section IV-B4). The optimal trajectory is selected based on maximum expected value under the IS distribution q :

$$E[V(\tau(\phi))] = \frac{1}{n_{IS}} \sum_{i=1}^{n_{IS}} \prod_{j=1}^N \frac{b(m_{i,j})}{q(m_{i,j})} V_i(\tau(\phi)), \quad (1)$$

$$\tau(\phi)^* = \arg \max E[V(\tau(\phi))],$$

where n_{IS} denotes the number of resampled scenarios, N is the number of exo-agents, ϕ is a sampled scenario, and $\tau(\phi)$ is a candidate trajectory generated under scenario ϕ . The trajectory’s value is represented by $V(\tau(\phi))$. The importance weights $\frac{b(m_{i,j})}{q(m_{i,j})}$ correct for sampling bias, ensuring unbiased estimation and thus robust trajectory selection.

VIII. EXPERIMENTS

We compare Hi-Drive with both learning-based and planning-based methods on a wide variety of real-world driving scenes, including those from the nuPlan benchmark and a private dataset. Results show that Hi-Drive generalized well across diverse urban environments, improving driving safety and ensuring successful task completion. Compared to learning-based methods, although Hi-Drive requires no training, it achieved state-of-the-art performance in reactive settings and competitive results in non-reactive settings on nuPlan. Compared to existing planners, Hi-Drive is the only planner that achieved zero collisions and 100% task success. Ablation studies emphasize the significant contributions of hierarchical POMDP planning and cross-scenario trajectory optimization. In the accompanying video, we also demonstrate how Hi-Drive’s reasoning process can be transparently interpolated.

A. Comparison with Learning-Based Models

We first compare Hi-Drive to state-of-the-art learning-based methods on the nuPlan benchmark [21]. The evaluation covers three datasets: *Val14* [34] (1,118 regular scenes), *Test14-random* [35] (268 random scenes), and *Test14-hard* [35] (272 complex scenes), with increasing difficulty. Each scenario lasts 15 seconds. We evaluate performance in both reactive and non-reactive settings using the official nuPlan metric, where higher scores indicate better performance.

Baselines include learning-based approaches—PLUTO [36] (contrastive imitation learning) and Diffusion Planner [37] (diffusion models)—as well as hybrid methods that refine learned trajectories using optimization, such as *PDM-Hybrid* [34] (winner of the 2023 nuPlan challenge), PLUTO w. refine, and Diffusion Planner w. refine.

Results in Table I demonstrate that Hi-Drive consistently outperforms learning-based and hybrid methods in reactive scenarios, and achieves competitive results in non-reactive scenarios. Notably, in challenging scenarios (*Test14-hard*), Hi-Drive significantly exceeds the performance of the best learning baseline. Note that Hi-Drive’s performance is achieved without training.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE I
COMPARISON WITH LEARNING-BASED MODELS.

Type	Planner	Val14		Test14-random		Test14-hard	
		R	NR	R	NR	R	NR
<i>Expert</i>	Log-replay	80.32	93.53	75.86	94.03	68.80	85.96
Learning-based	PLUTO	78.11	88.89	78.62	89.90	59.74	70.03
	Diffusion Planner	82.80	89.87	82.93	89.19	69.22	75.99
Hybrid	PDM-Hybrid	92.11	92.77	91.28	90.10	76.07	65.99
	PLUTO w/ refine.	76.88	92.88	90.29	92.23	76.88	80.08
	Diffusion Planner w/ refine.	<u>92.90</u>	94.26	91.75	94.80	82.00	78.87
Model-based	Hi-Drive (Ours)	93.15	<u>93.62</u>	<u>92.31</u>	93.71	83.18	81.41
	Hi-Drive-Dynamic (Ours)	93.15	93.51	92.56	<u>93.72</u>	<u>83.17</u>	<u>81.88</u>
	Hi-Drive-QCNet (Ours)	92.84	93.42	91.80	93.37	82.69	81.91
	Hi-Drive-Uniform (Ours)	92.80	92.60	91.62	92.08	82.56	79.98

TABLE II
COMPARISON WITH MODEL-BASED PLANNING.

Type	Planner	Private Long-Term					Val14		Test14-random		Test14-hard	
		NR					R	NR	R	NR	R	NR
		Reward \uparrow	Coll.R. (%) \downarrow	MGR (%) \downarrow	TTG (s) \downarrow	Comfort \uparrow	Score	Score	Score	Score	Score	Score
w/o unc.	Decoupled	-0.90	<u>1.70</u>	0.00	42.74	<u>0.96</u>	88.38	88.38	89.09	91.06	78.66	75.43
	ML-MCTS	<u>-0.88</u>	2.00	0.00	42.75	0.98	<u>93.01</u>	92.91	<u>92.08</u>	93.07	82.64	<u>81.30</u>
w/ unc.	Traj. Opt.	-2.43	10.00	0.90	36.46	0.92	72.63	72.99	70.32	71.95	65.00	60.75
	MPDM	-0.94	1.90	0.00	42.37	0.94	89.71	89.06	88.66	89.86	77.97	75.51
	POMDP	-0.89	1.80	<u>0.10</u>	<u>41.36</u>	0.95	92.92	<u>93.20</u>	91.80	<u>93.17</u>	<u>82.94</u>	80.70
	Hi-Drive (Ours)	-0.77	0.00	0.00	41.44	0.98	93.15	93.62	92.31	93.71	83.18	81.41

TABLE III
ABLATION STUDY RESULTS

Planner	Private Long-Term					Val14		Test14-random		Test14-hard	
	NR					R	NR	R	NR	R	NR
	Reward \uparrow	Coll.R. (%) \downarrow	MGR (%) \downarrow	TTG (s) \downarrow	Comfort \uparrow	Score	Score	Score	Score	Score	Score
W/O-Unc.	-0.88	2.00	0.00	42.75	0.98	<u>93.01</u>	92.91	92.08	93.07	82.64	<u>81.30</u>
W/O-CSTO	-0.89	1.80	<u>0.10</u>	41.36	<u>0.95</u>	92.92	93.20	91.80	93.17	<u>82.94</u>	80.70
W/O-IS	<u>-0.79</u>	<u>1.00</u>	<u>0.10</u>	<u>41.39</u>	0.98	92.92	<u>93.23</u>	92.09	93.42	82.88	81.26
H=4s	-0.81	0.00	1.10	43.09	0.98	92.98	93.02	<u>92.20</u>	<u>93.65</u>	82.66	80.43
H=9s (Full)	-0.77	0.00	0.00	41.44	0.98	93.15	93.62	92.31	93.71	83.18	81.41

* **Bold** indicates best; underscored indicates second-best.

We also evaluated several variants of Hi-Drive: *Hi-Drive-Uniform* maintains a fixed uniform belief over others' intentions, which results in overly conservative behaviors, showing the benefits of belief tracking; *Hi-Drive-Dynamic* considers the intention dynamics of exo-agents by sampling future intentions from a prior distribution. Its similar performance shows that dynamics of future intentions play an insignificant role for online planning; lastly, *Hi-Drive-QCNet* integrates trajectories predicted by a neural network (QCNet [38]), as an additional modality, to cover unstructured behaviors. Its similar performance shows that our selected driver model suite are indeed sufficient to model versatile behaviors.

B. Comparison with Model-Based Planning

We further evaluate Hi-Drive against model-based planners. In addition to nuPlan, we introduce the *Private Long-term* dataset, featuring 100 real-world scenes with 40-second episodes.

To reveal detailed planning performance, we introduce new metrics including the overall performance represented by the average cumulative reward (*reward*), safety performance represented by the collision rate among test episodes (*Coll.R.*), task completion performance characterized by the rate of missing the goal lane segment (*MGR*), driving efficiency denoted by the average time to reach goal (*TTG*), and a comfort score related to lateral acceleration (*Comfort*).

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

The compared planners include: *Decoupled*, which follows a given global route with IDM, *ML-MCTS* [17]–[20], which performs MCTS with maximum-likelihood scenarios, *Traj. Opt.* [16], which performs trajectory optimization under uncertainty, *MPDM* [2]–[4], representing multipolicy decision-making, and *POMDP* [6], [7], representing belief tree search planning without trajectory optimization.

Table II shows that Hi-Drive outperforms all planners across both benchmarks. On nuPlan, which requires adapting to diverse urban environments, Hi-Drive consistently ranks highest in reactive and non-reactive settings. On *Private Long-term*, which requires long-term planning and precise goal reaching, Hi-Drive achieves the safest driving and the highest task success.

Detailed analyses on *Private Long-term* highlight Hi-Drive’s advantages. Unlike planners without uncertainty modeling (*Decoupled* and *ML-MCTS*), Hi-Drive proactively mitigates risks, effectively reducing the collision rate to zero. Compared to planners with uncertainty modeling (*Traj. Opt.*, *MPDM*, *POMDP*), Hi-Drive effectively integrates behavior and trajectory planning, achieving the highest cumulative reward as well as a notable 0% miss-goal rate, while ensuring safety.

C. Ablation Study

In this section, we conduct ablation studies to evaluate key components of Hi-Drive, by comparing with 4 variants: *W/O-Unc.*, that uses a single maximum-likelihood scenario for planning; *W/O-CSTO*, which omits cross-scenario trajectory optimization; *W/O-IS*, which omits importance sampling in trajectory optimization; and *H=4s*, with shortened planning horizon from 9s to 4s. Results in Table III illustrate the contributions of each component. Removing POMDP planning (*W/O-Unc.*) increases collision from 0% to 2%, due to passive and late response to risks; Removing cross-scenario trajectory optimization (*W/O-CSTO*) also increases collision by 1.8%, due to insufficient granularity in trajectory analysis. Removing importance sampling (*W/O-IS*) in trajectory optimization reduces safety and goal-reaching, due to a lack of focus on rare and critical events. Lastly, shortening the planning horizon (*H=4s*) increases missed goals (1.1%), highlighting the necessity of long-term planning.

D. Computation Time

Hi-Drive achieves real-time performance when running in a single CPU thread, handling over 20 exo-agents on average. The planning time meets the typical 10 Hz requirement for AD systems, measuring 94.3 ms on an Intel i7-8700K, 91.6 ms on a Xeon Gold 6330, and 112.4 ms on a Xeon Gold 6240C CPU. For the i7-8700K test, time is mainly spent on belief tree search (45.0 ms), followed by trajectory optimization (21.1 ms) and belief tracking (13.6 ms).

E. Interpretable Reasoning

Fig. 1(b) and (c) illustrate the visual interpretation of the reasoning process of Hi-Drive in a scene of a right turn. A reduced number of driving scenarios are visualized for a clear

presentation. Vehicles in each scenario are denoted in a unique color. In belief tree search (Fig. 1b), a suboptimal action sequence (red) suggested turning into the rightmost lane, getting the ego-vehicle stuck in the traffic flow. In contrast, the optimal action sequence (green) suggests turning into the left lane, which is less crowded, thus enables efficient progress. During trajectory optimization (Fig. 1c), scenarios are resampled from an importance sampling distribution emphasizing critical agents (highlighted in red). Trajectories generated from the optimal action sequence are cross-evaluated with scenarios, selecting the green trajectory, which integrates seamlessly into traffic, and abandoned the red trajectory, which made little progress. More detailed visualization of the reasoning process is included in the accompanying video.

IX. CONCLUSION

This paper presents Hi-Drive, a hierarchical POMDP-based planner designed to handle uncertainties in behavior and trajectory planning for urban autonomous driving. Hi-Drive explicitly models uncertainties in other vehicles’ behavioral intentions and driving styles, allowing effective belief tracking and robust decision-making. Our method combines hierarchical behavior planning with importance sampling-based trajectory optimization, achieving safe, efficient, and comfortable driving in diverse and complex urban environments. Experiments on real-world driving datasets validate that Hi-Drive outperforms both learning-based and model-based planning approaches, demonstrating strong safety and robustness.

Hi-Drive still has the following limitations, inducing important directions of future work. First, the current set of driver models we adopted were designed or trained for urban environments with HD-maps available. Future work could extend this to map-less environments by adding map-prediction modules [39] or training new prediction models with map-less data. Second, we used a hand-crafted importance sampling distribution, which could be insufficient when handling very complex and crowded interactions; we plan to extend this to learned importance sampling distributions, for which [11] offers a potential approach. Furthermore, our current planner can be integrated with heuristic learning [40] and macro-action learning [41] to further enhance planning efficiency and performance. Parallelization can also be a significant aid to real-time performance, for example, through GPU parallelization [42] or SIMD parallelization via instruction sets [43].

ACKNOWLEDGMENT

Following the RAS policy, the use of generative AI is strictly limited to improve self-written texts to enhance readability. None of the presented methods and results (figures, equations, numbers, etc.) are generated by AI.

REFERENCES

- [1] J. Pineau, G. Gordon, and S. Thrun, “Anytime point-based approximations for large POMDPs,” *Journal of Artificial Intelligence Research*, vol. 27, pp. 335–380, 2006.
- [2] E. Galceran, A. G. Cunningham, R. M. Eustice, and E. Olson, “Multipolicy decision-making for autonomous driving via changepoint-based behavior prediction: Theory and experiment,” *Autonomous Robots*, vol. 41, no. 6, pp. 1367–1382, 2017.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

- [3] E. Schmerling, K. Leung, W. Vollprecht, and M. Pavone, "Multimodal probabilistic model-based planning for human-robot interaction," in *2018 IEEE International Conference on Robotics and Automation*, 2018, pp. 3399–3406.
- [4] W. Ding, L. Zhang, J. Chen, and S. Shen, "Epsilon: An efficient planning system for automated vehicles in highly interactive environments," *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 1118–1138, 2021.
- [5] S. Brechtel, T. Gindele, and R. Dillmann, "Probabilistic decision-making under uncertainty for autonomous driving using continuous POMDPs," in *IEEE Intelligent Transportation Systems Conference*, 2014, pp. 392–399.
- [6] Z. Sunberg and M. J. Kochenderfer, "Improving automated driving through POMDP planning with human internal states," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 11, pp. 20073–20083, 2022.
- [7] D. S. González, M. Garzón, J. S. Dibangoye, and C. Laugier, "Human-like decision-making for automated driving in highways," in *IEEE Intelligent Transportation Systems Conference*, 2019, pp. 2087–2094.
- [8] C. Xia, M. Xing, and S. He, "Interactive planning for autonomous driving in intersection scenarios without traffic signs," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 12, pp. 24818–24828, 2022.
- [9] K. Shu, H. Yu, X. Chen, S. Li, L. Chen, Q. Wang, L. Li, and D. Cao, "Autonomous driving at intersections: A behavior-oriented critical-turning-point approach for decision making," *IEEE/ASME Transactions on Mechatronics*, vol. 27, no. 1, pp. 234–244, 2021.
- [10] H. Bai, S. Cai, N. Ye, D. Hsu, and W. S. Lee, "Intention-aware online POMDP planning for autonomous driving in a crowd," in *2015 IEEE International Conference on Robotics and Automation*, 2015, pp. 454–460.
- [11] M. H. Danesh, P. Cai, and D. Hsu, "LEADER: Learning Attention over Driving Behaviors for Planning under Uncertainty," in *Conference on Robot Learning*. PMLR, 2023, pp. 199–211.
- [12] M. Naghshvar, A. K. Sadek, and A. J. Wiggers, "Risk-averse behavior planning for autonomous driving under uncertainty," *arXiv preprint arXiv:1812.01254*, 2018.
- [13] O. Speidel, J. Ruof, and K. Dietmayer, "Graph-based motion planning for automated vehicles using multi-model branching and admissible heuristics," in *2021 IEEE International Conference on Autonomous Systems*, 2021, pp. 1–5.
- [14] W. Xu, J. Pan, J. Wei, and J. M. Dolan, "Motion planning under uncertainty for on-road autonomous driving," in *2014 IEEE International Conference on Robotics and Automation*, 2014, pp. 2507–2512.
- [15] L. Lützw, Y. Meng, A. C. Armijos, and C. Fan, "Density planner: Minimizing collision risk in motion planning with dynamic obstacles using density-based reachability," in *2023 IEEE International Conference on Robotics and Automation*, 2023, pp. 7886–7893.
- [16] X. Li, Z. Sun, D. Cao, Z. He, and Q. Zhu, "Real-time trajectory planning for autonomous urban driving: Framework, algorithms, and verifications," *IEEE/ASME Transactions on Mechatronics*, vol. 21, no. 2, pp. 740–753, 2015.
- [17] C. Li, T. Trinh, L. Wang, C. Liu, M. Tomizuka, and W. Zhan, "Efficient Game-Theoretic Planning With Prediction Heuristic for Socially-Compliant Autonomous Driving," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10248–10255, 2022.
- [18] C.-J. Hoel, K. Driggs-Campbell, K. Wolff, L. Laine, and M. J. Kochenderfer, "Combining planning and deep reinforcement learning in tactical decision making for autonomous driving," *IEEE Transactions on Intelligent Vehicles*, vol. 5, no. 2, pp. 294–305, 2019.
- [19] D. Lenz, T. Kessler, and A. Knoll, "Tactical cooperative planning for autonomous highway driving using Monte-Carlo Tree Search," in *IEEE Intelligent Vehicles Symposium*, 2016, pp. 447–453.
- [20] S. V. Albrecht, C. Brewitt, J. Wilhelm, B. Gjevvar, F. Eiras, M. Dobre, and S. Ramamoorthy, "Interpretable goal-based prediction and planning for autonomous driving," in *2021 IEEE International Conference on Robotics and Automation*, 2021, pp. 1043–1049.
- [21] 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*, 2022.
- [22] H. Bey, M. Sackmann, A. Lange, and J. Thielecke, "Handling Prediction Model Errors in Planning for Automated Driving Using POMDPs," in *2021 IEEE International Intelligent Transportation Systems Conference*, Sep. 2021, pp. 439–446.
- [23] M. S. Saleem, R. Veerapaneni, and M. Likhachev, "A pomdp-based hierarchical planning framework for manipulation under pose uncertainty," *arXiv preprint arXiv:2409.18775*, 2024.
- [24] W. Zhao and W. Chen, "Hierarchical pomdp planning for object manipulation in clutter," *Robotics and Autonomous Systems*, vol. 139, p. 103736, 2021.
- [25] A. Jain and S. Niekum, "Efficient hierarchical robot motion planning under uncertainty and hybrid dynamics," in *Conference on Robot Learning*. PMLR, 2018, pp. 757–766.
- [26] S. C. Ong, S. W. Png, D. Hsu, and W. S. Lee, "Pomdp for robotic tasks with mixed observability," in *Robotics: Science and systems*, vol. 5, no. 4, 2009.
- [27] S. Sheng, D. Parker, and L. Feng, "Safe pomdp online planning via shielding," in *2024 IEEE International Conference on Robotics and Automation*. IEEE, 2024, pp. 126–132.
- [28] S. Carr, N. Jansen, S. Bharadwaj, M. T. Spaan, and U. Topcu, "Safe policies for factored partially observable stochastic games," 2021.
- [29] A. Kesting, M. Treiber, and D. Helbing, "General Lane-Changing Model MOBIL for Car-Following Models," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 1999, no. 1, pp. 86–94, 2007.
- [30] R. C. Coulter, *Implementation of the pure pursuit path tracking algorithm*. Carnegie Mellon University, The Robotics Institute, 1992.
- [31] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical Review E*, vol. 62, no. 2, pp. 1805–1824, 2000.
- [32] P. M. Djuric, J. H. Kotecha, J. Zhang, Y. Huang, T. Ghirmai, M. F. Bugallo, and J. Miguez, "Particle filtering," *IEEE Signal Processing Magazine*, vol. 20, no. 5, pp. 19–38, 2003.
- [33] N. Ye, A. Somani, D. Hsu, and W. S. Lee, "DESPOT: Online POMDP planning with regularization," *Journal of Artificial Intelligence Research*, vol. 58, pp. 231–266, 2017.
- [34] D. Dauner, M. Hallgarten, A. Geiger, and K. Chitta, "Parting with misconceptions about learning-based vehicle motion planning," in *Conference on Robot Learning*, 2023.
- [35] 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*. IEEE, 2024, pp. 14123–14130.
- [36] J. Cheng, Y. Chen, and Q. Chen, "PLUTO: Pushing the Limit of Imitation Learning-based Planning for Autonomous Driving," Apr. 2024, arXiv:2404.14327.
- [37] 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 *The Thirteenth International Conference on Learning Representations*, 2025.
- [38] Z. Zhou, J. Wang, Y.-H. Li, and Y.-K. Huang, "Query-centric trajectory prediction," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 17863–17873.
- [39] X. Xiong, Y. Liu, T. Yuan, Y. Wang, Y. Wang, and H. Zhao, "Neural map prior for autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 17535–17544.
- [40] P. Cai, Y. Luo, A. Saxena, D. Hsu, and W. S. Lee, "Lets-drive: Driving in a crowd by learning from tree search," in *Proceedings of Robotics: Science and Systems*, Freiburg/Breisgau, Germany, June 2019.
- [41] Y. Lee, P. Cai, and D. Hsu, "MAGIC: Learning Macro-Actions for Online POMDP Planning," in *Proceedings of Robotics: Science and Systems*, Virtual, July 2021.
- [42] P. Cai, Y. Luo, D. Hsu, and W. S. Lee, "Hyp-despot: A hybrid parallel algorithm for online planning under uncertainty," *The International Journal of Robotics Research*, vol. 40, no. 2-3, pp. 558–573, 2021.
- [43] W. Thomason, Z. Kingston, and L. E. Kavvaki, "Motions in microseconds via vectorized sampling-based planning," in *2024 IEEE International Conference on Robotics and Automation*. IEEE, 2024, pp. 8749–8756.