

# Multimodal Belief-Space Covariance Steering with Active Probing and Influence for Interactive Driving

Devodita Chakravarty<sup>1</sup>, John Dolan<sup>2</sup>, Yiwei Lyu<sup>3</sup>

**Abstract**—Autonomous driving in complex traffic requires reasoning under uncertainty. Common approaches rely on prediction-based planning or risk-aware control, but these are typically treated in isolation, limiting their ability to capture the coupled nature of action and inference in interactive settings. This gap becomes especially critical in uncertain scenarios, where simply reacting to predictions can lead to unsafe maneuvers or overly conservative behavior. Our central insight is that safe interaction requires not only estimating human behavior but also shaping it when ambiguity poses risks. To this end, we introduce a hierarchical belief model that structures human behavior across coarse discrete intents and fine motion modes, updated via Bayesian inference for interpretable multi-resolution reasoning. On top of this, we develop an active probing strategy that identifies when multimodal ambiguity in human predictions may compromise safety and plans disambiguating actions that both reveal intent and gently steer human decisions toward safer outcomes. Finally, a runtime risk-evaluation layer based on Conditional Value-at-Risk (CVaR) ensures that all probing actions remain within human risk tolerance during influence. Our simulations in lane-merging and unsignaled intersection scenarios demonstrate that our approach achieves higher success rates and shorter completion times compared to existing methods. These results highlight the benefit of coupling belief inference, probing, and risk monitoring, yielding a principled and interpretable framework for planning under uncertainty.

## I. INTRODUCTION

Reasoning under uncertainty is essential for autonomous driving in complex traffic. Advances in probabilistic human behavior prediction and risk-aware control have enabled autonomous vehicles (AVs) to capture multimodal behaviors, estimate intent, and manage state distributions for both safety and performance [1]–[4]. However, prediction and control are often treated as separate, sequential modules, overlooking the fact that an AV’s actions can influence human drivers’ behaviors. In interactive traffic, where decisions are co-dependent, this lack of feedback is problematic [5], [6].

Most existing methods remain reactive: they plan based on anticipated trajectories without actively reducing uncertainty or shaping human responses. While such strategies may suffice under low uncertainty, they become risky in multimodal situations where multiple futures are equally plausible [7],

[8]. In these cases, the AV is left with either unsafe single-mode plans or overly conservative maneuvers that hedge against all possibilities [9], [10]. Treating uncertainty as a fixed external input, rather than as something shaped by interaction, leaves AVs vulnerable to indecision, inefficiency, and elevated risk in safety-critical situations.

We argue that safe and efficient interaction in uncertain traffic requires more than passive prediction: it demands active shaping of human behavior when ambiguity poses risks. Our key idea is to couple inference and action by equipping AVs with probing strategies that both reduce uncertainty and gently bias human decisions toward cooperative outcomes.

The **main contributions** of this paper are threefold: (1) we introduce a hierarchical belief model that captures human driving behavior across discrete intents and motion-level modes, updated via Bayesian inference to enable coarse-to-fine reasoning about multimodal interactions, providing interpretable high-level insights while retaining motion-level precision, (2) we explicitly account for the interplay between prediction and action by proposing probing strategies for the robot that not only reduce ambiguity but also strategically influence human decisions toward safer outcomes, and (3) we design a runtime risk-aware layer based on Conditional Value-at-Risk (CVaR) that complements uncertainty quantification by safeguarding against adverse tail events, ensuring probing actions remain safe and tolerable. Together, these components yield a purposeful interaction framework, with improved safety and enhanced efficiency demonstrated in scenarios like lane-merging and unsignaled intersections when compared to state-of-the-art baselines.

## II. RELATED WORK

Predicting human behavior is fundamental to safe interaction in mixed-autonomy traffic. Unlike physical systems with unimodal dynamics, human driving decisions are inherently multimodal due to intent ambiguity (e.g., simultaneously considering yielding or merging aggressively). Capturing this uncertainty requires models that reason across both discrete intent and continuous motion levels. Bayesian inference provides effective online updates of such beliefs [11], and hierarchical structures have been proposed to couple these levels [12], [13]. This coarse-to-fine representation enables prediction of diverse futures while maintaining interpretability for decision-making. However, most prior work focuses on prediction alone, leaving open how to act effectively under multimodal beliefs.

Risk-aware planning frameworks address this by explicitly accounting for control uncertainty. Approaches include en-

<sup>1</sup>Devodita Chakravarty is with the Department of Mechanical Engineering, Indian Institute of Technology Kharagpur, India. devoditac@kgpian.iitkgp.ac.in

<sup>2</sup>John Dolan is with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA. jdolan@andrew.cmu.edu

<sup>3</sup>Yiwei Lyu is with the Department of Computer Science and Engineering, Texas A&M University, College Station, TX, USA. yiwei1lyu@tamu.edu

forcing hard risk bounds, adding risk penalties to objectives, using chance constraints under probabilistic models, and adopting CVaR, which is distribution-free and penalizes adverse tail events [14], [15]. Extensions to multimodal settings employ Gaussian mixtures [9] or game-theoretic formulations [16], building on theoretical foundations in [17]. While effective in managing risk, these methods remain reactive: they assume fixed uncertainty distributions and do not actively reduce ambiguity. This becomes especially problematic when competing hypotheses have similar likelihoods, leaving the system indecisive at precisely the most safety-critical moments.

Belief-space planning further incorporates uncertainty into the state representation. Covariance steering offers analytical formulations for controlling both mean and uncertainty evolution with closed-form solutions for linear Gaussian systems [18]. Extensions such as trajectory distribution control [19] and constrained formulations [20], [21] improve applicability. Belief-based human modeling with hierarchical intent–motion structures has also been explored [12], [13]. Yet these approaches primarily manage and propagate uncertainty rather than actively probing to reshape interaction.

The shift from passive observation to active information gathering has further shaped human-robot interaction. Early work in active perception focused on sensor control and information-theoretic planning [22]. More recently, probing has been applied to interaction: balancing exploration and exploitation under dual control [6], uncovering hidden parameters [8], or influencing latent states through strategic actions [5]. Probing has also been integrated with multimodal prediction [7] and studied in adversarial contexts using game-theoretic models [10]. Despite these advances, most approaches neglect risk-awareness and rely on flat (non-hierarchical) belief representations, limiting their suitability for safety-critical driving.

Closest to our work, [8] maximizes probing for behavior influence but relies on flat intent beliefs and lacks hierarchical modeling or risk-bounded probing. [9] applies CVaR-based risk-aware control under multimodal uncertainty but treats risk separately from intent inference, without actively reducing ambiguity. [7] integrates probing with multimodal prediction but remains primarily reactive, without hierarchical intent structures or runtime risk monitoring. In contrast, our approach unifies hierarchical intent–motion belief modeling, active probing that both disambiguates uncertainty and influences human behavior, and CVaR-based runtime risk adjustment. This combination enables safe, interpretable, and efficient decision-making in uncertain interactive driving scenarios.

### III. METHODOLOGY

Following our motivation, we aim to answer the following research questions:

- 1) How should autonomous vehicles represent and update their beliefs about human intent and motion in uncertain, multimodal intents?
- 2) When ambiguity in human behavior predictions arises, how can it be identified and how can strategies be

designed to actively disambiguate intent, so that autonomous vehicles can make safer and more informed decisions proactively rather than merely reacting?

- 3) Consider risk in interactive decision-making: what is a more effective way to quantify and manage this risk without sacrificing efficiency, so that autonomous decisions remain safe and responsive in uncertain, dynamic traffic scenarios?

To address Q1, Sec. III-A introduces a hierarchical belief model for human behavior prediction. For Q2, Sec. III-B presents active probing to reveal hidden information. For Q3, Sec. III-C details risk quantification from multimodal belief, while Sec. III-D develops covariance steering to actively shape human behavior and its associated risk. Finally, Sec. III-E brings these elements together through a bidirectional feedback loop.

#### A. Human Behavior Prediction with Hierarchical Belief Model

Let the ego agent’s state at time  $t$  be denoted by  $x(t) \in \mathbb{R}^n$ , and the observation of a surrounding agent be  $z(t)$ . We model the latent intent of the surrounding agent as a discrete random variable  $i \in \{1, \dots, I\}$ . These **intents** represent high-level, mutually exclusive behavioral goals, such as *yielding*, *merging aggressively*, or *cruising neutrally*. Our belief over these intents is captured by the probability vector  $\pi(t) = [\pi_1(t), \dots, \pi_I(t)]^T$ .

This high-level representation, however, is too coarse to capture the nuances of agent behavior. To address this, each intent  $i$  is further refined by a set of behavioral **modes**  $k \in \{1, \dots, K_i\}$ , with associated weights  $w_{i,k}(t)$ . These modes represent different ways an intent can be realized; for example, a *yielding* intent can be executed via a *gradual deceleration* mode or a *sudden braking* mode. Each mode  $(i, k)$  is formally defined by a target Gaussian distribution  $\mathcal{N}(\mu_{i,k}, \Sigma_{i,k})$ , where  $\mu_{i,k}$  encodes the nominal state that behavior aims to achieve and  $\Sigma_{i,k}$  captures the variability around it. These parameters are specified offline (e.g., via clustering of demonstrations or expert priors [23]) and remain fixed during online inference. These Gaussian parameters serve two roles: they define the likelihood model used in the Bayesian belief update (Eq. 1–3) and specify the target mean and covariance for the covariance-steering controller (Eq. 7).

To address Q1, we maintain a hierarchical belief over discrete intents and their associated motion modes. The process begins at the lowest level: given a new observation  $z(t)$ , we compute the likelihood of each mode  $(i, k)$ :  $\ell_{i,k}(t) = p(z(t) | i, k)$ , which serves as the atomic piece of evidence grounding our inference in the most recent sensor data.

With evidence for each individual mode, we ascend the hierarchy to the intent level. The aggregated likelihood for intent  $i$  is a weighted sum over its constituent modes:

$$\mathcal{L}[i] = \sum_{k=1}^{K_i} w_{i,k}(t) \ell_{i,k}(t). \quad (1)$$

---

**Algorithm 1** Hierarchical Belief Update
 

---

```

1: Input: Observation  $\mathbf{z}(t)$ , prior beliefs  $\pi(t), \{w_{i,k}(t)\}$ 
2: Output: Updated beliefs  $\pi(t+1), \{w_{i,k}(t+1)\}$ 
3: for  $i = 1$  to  $I$  do
4:   for  $k = 1$  to  $K_i$  do
5:      $\ell_{i,k}(t) \leftarrow p(\mathbf{z}(t) | i, k)$ 
6:   end for
7: end for
8: for  $i = 1$  to  $I$  do
9:    $\mathcal{L}[i] \leftarrow \sum_{k=1}^{K_i} w_{i,k}(t) \ell_{i,k}(t)$ 
10: end for
11:  $Z \leftarrow \sum_{j=1}^I \pi_j(t) \mathcal{L}[j]$ 
12: for  $i = 1$  to  $I$  do
13:    $\pi_i(t+1) \leftarrow \frac{\pi_i(t) \mathcal{L}[i]}{Z}$ 
14:    $Z_i \leftarrow \sum_{j=1}^{K_i} w_{i,j}(t) \ell_{i,j}(t)$ 
15:   for  $k = 1$  to  $K_i$  do
16:      $w_{i,k}(t+1) \leftarrow \frac{w_{i,k}(t) \ell_{i,k}(t)}{Z_i}$ 
17:   end for
18: end for
19: return  $\pi(t+1), \{w_{i,k}(t+1)\}$ 

```

---

This ensures that intent likelihood reflects the entire distribution of associated behaviors, not just the single best-fitting mode. Modes that were already considered probable ( $w_{i,k}(t)$  is high) contribute more significantly, creating natural temporal smoothing of the belief.

To combine aggregated likelihoods with our prior  $\pi(t)$ , we apply Bayes' rule. The normalization constant  $Z = \sum_{j=1}^I \pi_j(t) \mathcal{L}[j]$  represents the total probability of observing  $\mathbf{z}(t)$  across all intents and modes. The posterior intent probability is then:

$$\pi_i(t+1) = \frac{\pi_i(t) \mathcal{L}[i]}{Z}, \quad \forall i \in \{1, \dots, I\}. \quad (2)$$

Intents consistent with the new observation have their probabilities amplified, while inconsistent ones are suppressed.

Finally, mode weights *within* each intent are updated analogously:

$$w_{i,k}(t+1) = \frac{w_{i,k}(t) \ell_{i,k}(t)}{\sum_{j=1}^{K_i} w_{i,j}(t) \ell_{i,j}(t)}. \quad (3)$$

This is effectively a separate Bayesian update performed conditionally for each intent  $i$ , redistributing probability mass toward modes that better explain  $\mathbf{z}(t)$ .

Together, these coupled updates (Eq. 2 and 3) sharpen beliefs at both levels simultaneously, preventing incoherence where a mode remains highly likely ( $w_{i,k}$  high) even though its parent intent has become negligible ( $\pi_i \approx 0$ ). The full procedure is summarized in Algorithm 1.

### B. Reveal Hidden Information with Active Probing

In uncertain traffic scenarios, ambiguity about human intent and motion can compromise both safety and efficiency. Addressing **Q2**, which asks how ambiguity can be identified and disambiguated, we introduce an **active probing** strategy that deliberately seeks information to clarify uncertainty. Instead of passively reacting to predictions, the ego vehicle selects actions that reduce the uncertainty in future beliefs by exploring ambiguous scenarios.

The probing objective is based on **entropy**, which measures uncertainty in the belief distributions over intents and motion modes. The entropy at time  $t$  is computed as:

$$H(\pi(t), w(t)) = - \sum_{i=1}^I \pi_i(t) \log \pi_i(t) - \sum_{i=1}^I \sum_{k=1}^{K_i} w_{i,k}(t) \log w_{i,k}(t) \quad (4)$$

where  $\pi_i(t)$  is the belief over intent  $i$  and  $w_{i,k}(t)$  is the belief over mode  $k$  within intent  $i$ .

To encourage actions that actively gather information, we define the probing objective as the expected reduction in uncertainty:

$$J_{\text{probe}}(u) = -\lambda_H \mathbb{E}[H(\pi(t+1), w(t+1))], \quad (5)$$

where  $\lambda_H > 0$  scales the importance of probing, and the expectation is taken over possible future observations that result from applying control  $u$ . By minimizing this objective, the ego vehicle selects actions that are expected to reveal hidden aspects of human behavior, making subsequent interaction safer and more informed.

### C. Uncertainty Quantification with Multimodal Belief

For **Q3**, which concerns how to effectively quantify and manage risk without sacrificing efficiency, we formalize the ego's uncertainty using a hierarchical belief structure. This structure explicitly represents the uncertainty across multiple behavioral modes and intents.

The belief over intents at time  $t$  is expressed as:  $\pi(t) = [\pi_1(t), \dots, \pi_I(t)]$ , where  $\pi_i(t)$  represents the probability of intent  $i$ . For each intent, the belief over motion modes is given by:  $w(t) = \{w_{i,k}(t)\}, i = 1, \dots, I, k = 1, \dots, K_i$ , where  $w_{i,k}(t)$  represents the probability of mode  $k$  for intent  $i$ . Combining these, the joint probability of the agent following mode  $(i, k)$  is:  $\omega_{i,k}(t) = \pi_i(t) w_{i,k}(t)$ , which reflects the agent's overall behavior under the updated belief.

To account for safety under deep uncertainty, we incorporate **Conditional Value-at-Risk (CVaR)**, which focuses on the worst-case scenarios rather than the average behavior. Given a set of  $S$  simulated future trajectories with associated costs  $J_{i,k}^{(s)}$ , the CVaR at confidence level  $\alpha$  is defined as:

$$\text{CVaR}_\alpha(i, k) = \frac{1}{m} \sum_{s=1}^m J_{i,k}^{(s)\uparrow}, \quad m = \max(1, \lceil \alpha S \rceil), \quad (6)$$

where  $J_{i,k}^{(s)\uparrow}$  are the sorted costs in descending order and  $\alpha \in (0, 1)$  is a chosen risk tolerance parameter. This approach allows the ego vehicle to balance both likelihood and potential danger, thereby quantifying uncertainty in a manner that explicitly integrates safety and efficiency considerations.

### D. Actively Shaping Human Belief via Covariance Steering

To actively shape human behavior while ensuring safety and efficiency, we address both **Q2** and **Q3** by formulating a finite-horizon covariance steering optimization problem over horizon  $T$ . This framework enables the ego vehicle to guide its own state trajectory toward desired outcomes

while controlling the associated uncertainty distribution by explicitly considering both expected behavior and worst-case risk. For each mode  $(i, k)$ , the target behavior, which is the velocity gaussian of the ego's belief on human intent, is specified by a Gaussian distribution  $\mathcal{N}(\mu_{i,k}, \Sigma_{i,k})$ , where  $\mu_{i,k}$  encodes the desired state and  $\Sigma_{i,k}$  encodes the allowable variability. These targets are determined offline through expert knowledge, clustering of observed trajectories, or prior experience [23] and reflect safe and cooperative behavior patterns. The ego's predicted state at time  $t$  is modeled by  $\mathcal{N}(\mu(t), \Sigma(t))$ , where  $\mu(t)$  and  $\Sigma(t)$  represent the predicted mean and covariance at time  $t$  given the applied control sequence.

The ego vehicle's objective is to minimize both the expected uncertainty and the deviation from desired behavior. This is compactly expressed using the previously defined objectives:

$$u^* = \operatorname{argmin}_u J(u) = J_{\text{probe}}(u) + J_{\text{influence}}(u), \quad (7)$$

where  $J_{\text{probe}}(u)$  captures the information gain through active probing (Eq. 5) and  $J_{\text{influence}}(u)$  quantifies alignment with the target behavior using the Kullback-Leibler divergence:

$$J_{\text{influence}}(u) = D_{\text{KL}}(\mathcal{N}(\mu(t), \Sigma(t)) \parallel \mathcal{N}(\mu_{i,k}, \Sigma_{i,k})). \quad (8)$$

The ego's state evolves according to discrete-time dynamics [24]:

$$\mu(t + \tau + 1) = A\mu(t + \tau) + Bu(t + \tau), \quad (9)$$

$$\Sigma(t + \tau + 1) = A\Sigma(t + \tau)A^\top + BWB^\top, \quad (10)$$

where  $A$  and  $B$  are linear approximations of the system dynamics, and  $W$  is the process noise covariance matrix.

To explicitly ensure robustness against worst-case scenarios, we impose Conditional Value-at-Risk (CVaR) constraints for each mode:

$$\text{CVaR}_\alpha(i, k) \leq \bar{J}, \quad (11)$$

where  $\bar{J}$  is a predefined risk threshold and  $\alpha \in (0, 1)$  is the confidence level for tail risk. The control inputs are constrained to remain within safe operational bounds:  $u_{\min} \leq u(t + \tau) \leq u_{\max}$ ,  $\forall \tau = 0, \dots, T - 1$ , ensuring that the vehicle operates within actuator limits and avoids unsafe maneuvers. State-of-the-art quadratic programming solvers such as OSQP (Operator Splitting Quadratic Program) can be used to solve this optimization problem. Even though this formulation offers a mathematically sound method with precise goals and limitations, it could be computationally costly to solve the entire optimization at each time step. Approximate techniques with warm-starts, like sequential quadratic programming or model predictive control (see Sec IV-A), can be used in real-world applications. Furthermore, solutions can be improved in between solver iterations using gradient-based updates, guaranteeing responsiveness in real-time interactions [25], [26]. By explicitly defining this covariance steering optimization, we create a principled framework that integrates safety, efficiency, and information gain while allowing for solver-based implementation that can be tailored to real-time operational requirements.

---

### Algorithm 2 Active Multimodal Covariance Steering with Risk-aware Probing

---

```

1: Input: Initial state  $x(0)$ , target distributions  $\{\mu_{i,k}, \Sigma_{i,k}\}$ , horizon  $T$ 
2: Output: Trajectory  $\{x(t)\}_{t=0}^T$ 
3: Initialize beliefs:  $\pi_i(0) \leftarrow 1/I$ ,  $w_{i,k}(0) \leftarrow 1/K_i$ 
4: for  $t = 0$  to  $T - 1$  do
5:   Observe environment:  $z(t)$ 
6:   Update beliefs:  $(\pi(t + 1), w(t + 1)) \leftarrow$ 
   BELIEFUPDATE( $z(t), \pi(t), w(t)$ )
7:   Compute weighted beliefs:  $\omega_{i,k} \leftarrow \pi_i(t + 1) w_{i,k}(t + 1)$ 
8:   Estimate CVaR:  $\text{CVaR}_\alpha(i, k)$  for each mode
9:    $u_{\text{total}} \leftarrow 0$ 
10:  for all  $(i, k)$  such that  $\omega_{i,k} > \epsilon$  do
11:    Formulate probing objective  $J_{\text{probe}}(u)$ 
12:    Formulate influence objective  $J_{\text{influence}}(u)$ 
13:     $J(u) \leftarrow J_{\text{probe}}(u) + J_{\text{influence}}(u)$ 
14:     $\nu_{i,k}(t) \leftarrow \operatorname{argmin}_u J(u)$ 
15:    subject to  $\{\mu_{i,k}, \Sigma_{i,k}\}$ ,
16:     $\text{CVaR}_\alpha(i, k) \leq \bar{J}$ ,
17:     $u_{\min} \leq u \leq u_{\max}$ 
18:     $u_{\text{total}} \leftarrow u_{\text{total}} + \omega_{i,k} \nu_{i,k}(t)$ 
19:  end for
20:  Apply control:  $x(t + 1) \leftarrow f(x(t), u_{\text{total}}, \Delta t)$ 
21:  if SafetyViolation( $x(t + 1)$ ) then
22:    break
23:  end if
24: end for
25: return  $\{x(t)\}_{t=0}^T$ 

```

---

#### E. Bi-Directional Prediction-Action Feedback Loop

A closely coupled feedback loop that incorporates belief updates, probing, and covariance steering makes it possible for prediction and action to interact. Every time step, the following sequence is carried out by the ego vehicle:

- 1) **Perceive:** The ego gathers a fresh observation  $z(t)$  from its sensors, capturing the motion of agents in the immediate vicinity.
- 2) **Infer:** The ego propagates uncertainty across both high-level intents and motion modes by using the Hierarchical Belief Update to refine its belief distributions  $\pi(t)$  and  $w_{i,k}(t)$  in light of the observation.
- 3) **Decide & Act:** By minimizing the objective  $J(u)$ , the probing controller integrates the influence and probing objectives to calculate the control action  $u(t)$ . This choice aligns the ego's trajectory towards safe and cooperative behaviors, takes into consideration the current uncertainty, and gives information acquisition priority in ambiguous situations.
- 4) **Propagate:** The ego's state is updated as:  $x(t + 1) \leftarrow f(x(t), u(t), \Delta t)$ , where the updated state depends on the applied control and time step.

Algorithm 2 summarizes this bi-directional feedback loop, where the ego's actions continuously reshape its beliefs, which in turn guide subsequent actions. The ego actively seeks information to clarify intent when entropy is high; when uncertainty decreases, it shifts to efficient execution, focusing on cooperative actions that are consistent with low-risk pathways. The ego balances exploration and exploitation in real time by dynamically modifying its trajectory and related uncertainty through covariance steering. The resulting interactions provide a practical and consequence-aware

approach to uncertainty-aware planning in complex traffic situations, consistent with human-driven behaviors, safer, and easier to interpret.

#### IV. SIMULATION & DISCUSSION

##### A. Experimental Setup

We evaluate the framework in two interactive driving scenarios: lane merging and an unsignaled four-way intersection, chosen for their challenges in longitudinal merging and lateral conflict resolution. Human-driven agents are modeled deterministically, reacting predictably to the ego vehicle’s probing. The human model follows [27], which generates best-response plans updated by risk thresholds and implicit communication. We adapt it as a reactive trajectory generator within our hierarchical belief structure, allowing integration with Bayesian updates and active probing so the ego can refine beliefs and influence human responses toward safe, cooperative outcomes.

Both the scenarios are run for 500 runs each. Initial positions and velocities of all vehicles are sampled uniformly. In lane merges, ego velocity is between 7–10 m/s, with surrounding traffic between 8–12 m/s. In intersections, ego velocity is 4–6 m/s and the crossing agent 5–8 m/s. Vehicle positions are randomized around the merge point or intersection center. Horizons are  $T = 4$  s (merge) and  $T = 6$  s (intersection) with a 10 Hz rate. Risk-sensitivity is incorporated via CVaR regularization, with  $\alpha \in [0.05, 0.2]$  to balance tail-risk and efficiency. Across runs, best-performing parameters were  $S = 100$ ,  $\alpha = 0.05$ ,  $\beta = 5$ ,  $\lambda_H = 0.5$ ,  $\gamma = 0.95$ ,  $\epsilon = 10^{-5}$ , horizon length of 30 steps ( $\Delta t = 0.1$ , corresponding to  $T = 3$  s), and  $K_i = 3$  intent modes. We solved the optimization by solving a Model Predictive Control (MPC) at each time step. Here, the problem is repeatedly solved over a finite horizon with updated state estimates, and only the first control input is applied before the optimization is resolved at the next time step.

##### B. Quantitative Analysis of Proposed Method

In order to evaluate the effectiveness of our proposed approach, we considered key performance metrics that reflect both safety and efficiency in interactive driving scenarios. These include *Success Rate*, which measures the percentage of simulation runs where the ego vehicle safely completes the maneuver; *Time to Merge* or *Time to Cross*, which indicates how quickly interactions are resolved; *Gap to Vehicle*, representing the distance maintained from surrounding agents; and *Velocity*, which reflects how efficiently the ego vehicle moves toward its objective. Additionally, *Longitudinal* and *Angular Jerk* are used to quantify the smoothness of maneuvers, ensuring that rapid or uncomfortable control changes are avoided. These metrics collectively capture how well the ego vehicle balances safety, efficiency, and comfort under uncertainty.

As shown in Table I and Table II (last columns), our method consistently outperforms the baseline approaches across both the lane merging and unsignaled intersection scenarios. For example, we achieved a success rate of **96%**

in lane merging and **94%** in intersections, demonstrating the robustness of our hierarchical belief model combined with risk-aware probing. The reduction in *Time to Merge* (to **2.03 ± 1.38 s**) and *Time to Cross* (to **4.14 ± 1.87 s**) shows that our method enables faster decision-making without compromising safety. The improvements in *Gap to Vehicle* and *Velocity* reflect how the ego vehicle can safely operate closer to other agents and at higher speeds, leveraging the CVaR-based tail-risk adjustment to ensure that such actions remain within acceptable safety bounds. Furthermore, the comparable or improved values for *Longitudinal* and *Angular Jerk* indicate that these efficiency gains do not result in abrupt or uncomfortable maneuvers. These results highlight how actively shaping uncertainty through probing and influence, rather than passively reacting, enables safer, smoother, and more efficient interactions in complex traffic environments.

##### C. Example Case Studies

Fig. 1 presents a step-by-step illustration of the autonomous vehicle (ego) interacting with a human-driven vehicle during a lane merging maneuver. The figure visualizes the progression of the hierarchical belief update, risk assessment, control effort, and resulting trajectories across time steps  $t_1$  through  $t_5$ . At  $t_1$  (Uncertain Initial Belief), the ego vehicle starts with a prior hierarchical belief over the human driver’s latent intent, distributed over the three discrete intent categories: aggressive, neutral, and cooperative. This belief is represented probabilistically as the vector  $\pi(t_1) = [\pi_{\text{aggressive}}, \pi_{\text{neutral}}, \pi_{\text{cooperative}}](t_1)$ , reflecting the initial uncertainty. Each intent further decomposes into multiple motion modes  $k$ , each with corresponding Gaussian targets  $\mathcal{N}(\mu_{i,k}, \Sigma_{i,k})$ . The prior belief is broad due to limited observation at this early stage, reflecting ambiguity about how the human will behave. At  $t_2$  (Active Probing Increases Aggressive Intent Likelihood), the ego executes an active probing maneuver designed to reduce uncertainty by influencing the human’s behavior and eliciting informative responses. This maneuver is computed by minimizing the expected future entropy of the hierarchical belief and steering the system to maximize information gain about human intent (Eq. 4, 5). As a result, after observing the human response to this probing action, the Bayesian belief update increases the probability weight on the aggressive intent mode,  $\pi_{\text{aggressive}}(t_2)$ . Simultaneously, the tail risk of collision or unsafe interaction rises, as measured by the Conditional Value-at-Risk (CVaR) metric, since aggressive behavior carries higher risk in the merge context. Consequently, the control effort by the ego vehicle increases at  $t_2$ , as it adapts its maneuvers to safely accommodate the possibility of aggressive human actions while maintaining efficiency constraints. At  $t_3$  (Human Response Shifts Belief Toward Neutral), the human driver reacts to the ego’s probing by adjusting their motion (velocity), which is observed and incorporated in the belief update. This results in shifting posterior belief from aggressive towards the neutral intent mode,  $\pi_{\text{neutral}}(t_3)$ , reflecting reduced risk. The CVaR risk

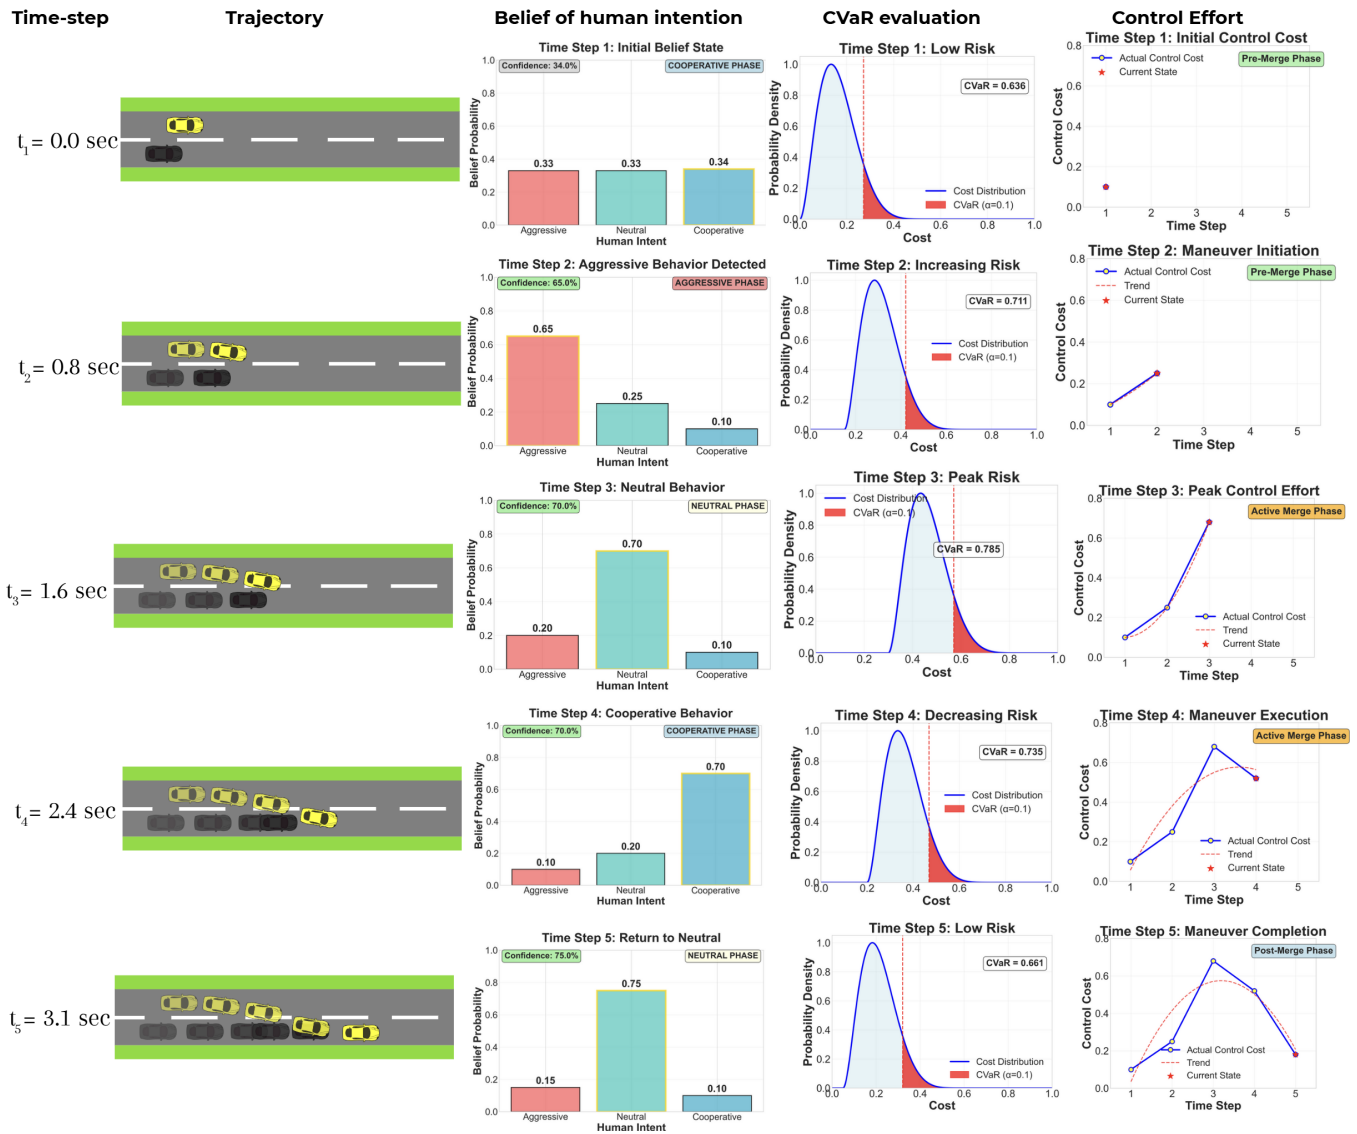


Fig. 1: Overview of the lane-merging interaction between an autonomous ego vehicle and a human-driven vehicle, illustrating the evolution of hierarchical belief updates, CVaR-based risk assessment, control effort, and resulting trajectories from  $t_1$  to  $t_5$ . The progression shows how active probing and Bayesian updates refine intent inference from uncertain to cooperative, while control effort adapts to balance safety and efficiency. Specifically, the *hierarchical belief update* displays Bayesian updates integrating noisy observations and probing-induced responses, evolving from a uniform intent distribution at  $t_1$  to a concentrated cooperative belief at  $t_5$ . Instead of passively observing, the *active probing controller* at  $t_2$  demonstrates how probing actions intentionally change the interaction and belief state to reveal latent human intentions. The *CVaR risk assessment* curve illustrates how risk-aware control adjusts the vehicle's aggressiveness and conservatism during the interaction by first rising and then falling over time steps. Lastly, the model's ability to dynamically combine efficiency and safety is demonstrated by the way the *adaptive control effort* changes in response to inferred risk and intent certainty.

estimation correspondingly decreases as the system's uncertainty clarifies and less adverse outcomes are forecasted. The ego's control effort adjusts downward, balancing safety with the new, safer predicted human behavior. At  $t_4$  (Human Yields; Cooperative Intent Becomes Most Likely), further interaction by probing and continuous Bayesian updates refines the belief and influences the human's intention to yield, reflected as increased probability on the cooperative intent mode. This cooperative mode corresponds to trajectories with smoother, less adversarial human motions that facilitate a safe merge. The reduction in CVaR risk continues, as the threat of collision or emergency braking diminishes. The ego vehicle exploits this clarity to optimize for smoother, more

efficient control inputs. At  $t_5$  (Maneuver Completion with Reduced Control Effort), the ego confidently completes the lane merge after the human's intent has been sufficiently resolved and cooperative behavior verified. There is little remaining uncertainty indicated by the low entropy of the hierarchical belief distribution, which is centred around the cooperative mode. A safe interaction envelope is reflected in the low CVaR-driven tail-risk measure. The lowest level of control efforts is a sign of effective, well-thought-out execution made possible by earlier active probing and risk-conscious planning.

The intersection scenario (Fig. 2) further highlights probing as signaling. Both vehicles arrive simultaneously with




Fig. 2: Unsignaled four-way intersection. The ego shapes the crossing agent’s belief, leading it to yield.

no right of way. The ego performs small, risk-bounded probes that are interpreted as a commitment to cross first, prompting the human to yield. This avoids deadlock and resolves conflict while keeping actions within CVaR limits.

#### D. Comparative Analysis with Baselines

We compare our approach against three representative baselines from the literature: AP-IH [8], which explicitly formulates an information-maximizing probing objective to reveal human-model parameters and then uses the gathered information to influence behavior; CC-MPC [9], a state-of-the-art trajectory planner that handles multimodal obstacle uncertainty using Gaussian mixture models and chance/CVaR approximations; and AP-MP [7], which couples multimodal prediction with an active probing strategy to reduce prediction ambiguity during planning. These baselines were chosen because they represent (i) active-probing formulations, (ii) robust risk-aware planning under multimodal uncertainty, and (iii) integrated probing and prediction approaches, i.e., they are the most relevant prior works to which our multimodal hierarchical, risk-aware probing framework can be fairly compared. To the best of our knowledge, none of these prior works combine a multimodal hierarchical intent-mode belief representation with an explicit tail-risk (CVaR) adjustment of multimodal weights and an entropy-minimizing probing objective; the three baselines above are therefore the closest existing methods.

Table I (lane merge) and Table II (unsignaled four-way intersection) summarize the quantitative comparison over 500 runs per planner. We observe our method outperforms consistently across all parameters. It is to be noted that the headway gap is less in our method, which is deliberate and safe; it reflects an efficiency-driven decision. Our CVaR-based evaluation explicitly penalizes high-tail-cost outcomes, so the controller reduces conservative spacing when the tail risk is acceptable, producing faster merges while keeping worst-case outcomes constrained. The baselines differ in performance based on how they address uncertainty, probing, and risk. AP-IH focuses on information gain through probing but ignores risk and intent inference, resulting in safe yet slow maneuvers with lower success rates under high uncertainty. CC-MPC improves by incorporating risk-aware planning with chance constraints and Gaussian mixtures, but its assumption of fixed uncertainty makes it overly cautious

and less adaptable. AP-MP further reduces ambiguity by combining prediction and probing, yet its flat intent model and lack of tail-risk handling leave it vulnerable to rare but critical failures. Our method surpasses these baselines through updating hierarchical beliefs, shaping interactions through probing and influence, and adjusting for tail risk with CVaR. This enables the ego vehicle to safely gather and act on information, achieving faster, more confident maneuvers without compromising safety. The results highlight the importance of coupling belief updates with action for effective navigation in uncertain interactive driving.

Fig. 3 illustrates qualitative trajectories for the same example trial lane merge scenario (with the same parameters) across all methods. We observe that AP-IH (brown) adopts a highly conservative approach, requiring longer time to merge; CC-MPC (pink) manages to merge somewhat faster but exhibits indecision and prolonged lateral alignment with human vehicles; and AP-MP (blue) achieves smoother motion but still requires more time to complete the merge. In contrast, our method (yellow) executes a confident and decisive maneuver, merging more quickly and with smoother transitions. This visualization aligns with the quantitative findings in Table I, highlighting how our probing-driven belief updates enable faster resolution of interactions, fewer deadlocks, and higher overall efficiency. While our approach achieves strong overall performance, we observe failures in certain edge cases such as dense traffic with very high traffic speed, which makes it difficult for the ego to take meaningful probe actions. In these situations, our probing actions may not sufficiently disambiguate aggressive human intent, causing the ego to mispredict behaviors and commit to unsafe trajectories, sometimes resulting in crashes or incomplete maneuvers. Therefore, future work includes exploring multi-step lookahead probing strategies to improve robustness.

## V. CONCLUSIONS

We propose a principled framework that establishes a bidirectional feedback loop between control and prediction in interactive driving. Specifically, we introduce a risk-aware multimodal covariance steering method that actively learns and shapes human behavior, modeled through a hierarchical belief framework for intent inference. By proactively reducing ambiguity in predicting human actions, our approach improves safety and efficiency in scenarios such as lane merging and unsignaled intersections.

## REFERENCES

- [1] L. Chen, J. Wang, T. Mortlock, P. Khargonekar, and M. A. Al Faruque, “Hyperdimensional uncertainty quantification for multimodal uncertainty fusion in autonomous vehicles perception,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 22306–22316, 2025.
- [2] E. Candela, O. Doustaly, L. Parada, F. Feng, Y. Demiris, and P. Angeloudis, “Risk-aware controller for autonomous vehicles using model-based collision prediction and reinforcement learning,” *Artificial Intelligence*, vol. 320, p. 103923, 2023.
- [3] D. Zheng, J. Ridderhof, P. Tsiotras, and A.-a. Agha-mohammadi, “Belief space planning: A covariance steering approach,” in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 11051–11057, IEEE, 2022.

TABLE I: Comparison of Performance Metrics for lane merging scenario for 500 simulation runs

Metric	AP-IH	CC-MPC	AP-MP	Ours
Success Rate	82%	87%	92%	<b>96%</b>
Time to Merge (s)	6.80 ± 1.34	4.72 ± 1.65	5.25 ± 1.42	<b>2.03 ± 1.38</b>
Gap to Vehicle (m)	9.54 ± 3.60	9.84 ± 2.10	8.10 ± 2.20	<b>7.60 ± 2.29</b>
Velocity (m/s)	8.35 ± 1.06	7.90 ± 1.08	8.01 ± 1.17	<b>8.88 ± 1.02</b>
Longitudinal Jerk (m/s <sup>3</sup> )	0.31 ± 0.18	0.28 ± 0.19	<b>0.21 ± 0.22</b>	0.22 ± 0.16
Angular Jerk (rad/s <sup>3</sup> )	(1.52 ± 1.21) × 10 <sup>-2</sup>	(1.84 ± 1.10) × 10 <sup>-2</sup>	(1.76 ± 1.05) × 10 <sup>-2</sup>	<b>(1.49 ± 0.98) × 10<sup>-2</sup></b>

TABLE II: Comparison of Performance Metrics for an unsignaled four-way intersection scenario for 500 simulation runs

Metric	AP-IH	CC-MPC	AP-MP	Ours
Success Rate	78%	84%	88%	<b>94%</b>
Time to Cross (s)	8.02 ± 2.04	5.78 ± 1.90	6.80 ± 2.90	<b>4.14 ± 1.87</b>
Gap to Vehicle (m)	6.20 ± 2.10	6.80 ± 2.00	6.10 ± 1.90	<b>5.70 ± 2.00</b>
Velocity (m/s)	4.56 ± 0.78	4.73 ± 0.72	4.84 ± 0.67	<b>5.11 ± 0.61</b>
Longitudinal Jerk (m/s <sup>3</sup> )	0.21 ± 0.15	0.20 ± 0.16	0.19 ± 0.14	<b>0.16 ± 0.13</b>
Angular Jerk (rad/s <sup>3</sup> )	0.18 ± 0.12	0.15 ± 0.11	<b>0.12 ± 0.10</b>	0.14 ± 0.11




Fig. 3: Trajectory comparison across methods for the same lane merge scenario. AP-IH (brown), CC-MPC (pink), and AP-MP (blue) manage to merge but require longer times or more conservative maneuvers, while our method (yellow) merges more quickly and with smoother motion, reflecting the efficiency and risk-aware confidence of our planner.

- [4] J. Pilipovsky and P. Tsiotras, “Data-driven covariance steering control design,” in *2023 62nd IEEE Conference on Decision and Control (CDC)*, pp. 2610–2615, IEEE, 2023.
- [5] D. Sadigh, S. S. Sastry, S. A. Seshia, and A. Dragan, “Information gathering actions over human internal state,” in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 66–73, IEEE, 2016.
- [6] H. Hu and J. F. Fisac, “Active uncertainty reduction for human-robot interaction: An implicit dual control approach,” in *International Workshop on the Algorithmic Foundations of Robotics*, pp. 385–401, Springer, 2022.
- [7] D. Gadginmath, F. Nawaz, M. Sung, F. M. Tariq, S. Bae, D. Isele, F. Pasqualetti, and J. D’sa, “Active probing with multimodal predictions for motion planning,” 2025.
- [8] S. Wang, Y. Lyu, and J. M. Dolan, “Active probing and influencing human behaviors via autonomous agents,” *arXiv preprint arXiv:2304.11817*, 2023.
- [9] K. Ren, H. Ahn, and M. Kamgarpour, “Chance-constrained trajectory planning with multimodal environmental uncertainty,” *IEEE Control Systems Letters*, vol. 7, pp. 13–18, 2022.
- [10] S. He, Y. Tao, I. Spasojevic, V. Kumar, and P. Chaudhari, “An active perception game for robust information gathering,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 14168–14174, IEEE, 2025.
- [11] E. I. Farhi and V. Indelman, “Bayesian incremental inference update by re-using calculations from belief space planning: a new paradigm,” *Autonomous Robots*, vol. 46, no. 7, pp. 783–816, 2022.
- [12] N. Ludlow, Y. Lyu, and J. Dolan, “Hierarchical learned risk-aware planning framework for human driving modeling,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2223–2229, IEEE, 2024.
- [13] L. Bao, Y. Pan, T. Peng, D. Kanoulas, and C. Zhou, “Hierarchical intention-aware expressive motion generation for humanoid robots,” *arXiv preprint arXiv:2506.01563*, 2025.
- [14] J. Yin, Z. Zhang, and P. Tsiotras, “Risk-aware model predictive path integral control using conditional value-at-risk,” *arXiv preprint arXiv:2209.12842*, 2022.
- [15] M. Ahmadi, X. Xiong, and A. D. Ames, “Risk-averse control via cvar barrier functions: Application to bipedal robot locomotion,” *IEEE Control Systems Letters*, vol. 6, pp. 878–883, 2021.
- [16] R. Chandra, M. Wang, M. Schwager, and D. Manocha, “Game-theoretic planning for autonomous driving among risk-aware human drivers,” in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 2876–2883, IEEE, 2022.
- [17] K. Ren, *Safe and risk-aware trajectory planning under Gaussian Mixture Model environmental uncertainty*. PhD thesis, University of British Columbia, 2022.
- [18] F. Liu, G. Rapakoulias, and P. Tsiotras, “Optimal covariance steering for discrete-time linear stochastic systems,” *IEEE Transactions on Automatic Control*, 2024.
- [19] J. Yin, Z. Zhang, E. Theodorou, and P. Tsiotras, “Trajectory distribution control for model predictive path integral control using covariance steering,” in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 1478–1484, IEEE, 2022.
- [20] I. M. Balci and E. Bakolas, “Constrained multi-modal density control of linear systems via covariance steering theory,” *arXiv preprint arXiv:2501.02866*, 2025.
- [21] I. M. Balci, E. Bakolas, B. Vlahov, and E. A. Theodorou, “Constrained covariance steering based tube-mpci,” in *2022 American Control Conference (ACC)*, pp. 4197–4202, IEEE, 2022.
- [22] A. Elfes, “Dynamic control of robot perception using stochastic spatial models,” *Information Processing in Mobile Robots, Springer-Verlag, Berlin*, 1991.
- [23] M. V. N. de Zepeda, F. Meng, J. Su, X.-J. Zeng, and Q. Wang, “Dynamic clustering analysis for driving styles identification,” *Engineering applications of artificial intelligence*, vol. 97, p. 104096, 2021.
- [24] K. Okamoto, “Optimal covariance steering: Theory and its application to autonomous driving,” *Ph. D. dissertation, Georgia Institute of Technology*, 2019.
- [25] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, “Osqp: An operator splitting solver for quadratic programs,” *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, 2020.
- [26] S. Hall, L. Ortmann, M. Picallo, and F. Dörfner, “Real-time projected gradient-based nonlinear model predictive control with an application to anesthesia control,” in *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 6193–6198, IEEE, 2022.
- [27] O. Siebinga, A. Zgonnikov, and D. A. Abbink, “A model of dyadic merging interactions explains human drivers’ behavior from control inputs to decisions,” *PNAS nexus*, vol. 3, no. 10, p. pgae420, 2024.