

Online Risk-Bounded Motion Planning for Autonomous Vehicles in Dynamic Environments

Xin Huang, Sungkweon Hong, Andreas Hofmann, Brian C. Williams

MIT Computer Science and Artificial Intelligence Laboratory

32 Vassar St., Cambridge, MA 02139

{xhuang, sk5050, hofma, williams}@csail.mit.edu

*

Abstract

A crucial challenge to efficient and robust motion planning for autonomous vehicles is understanding the intentions of the surrounding agents. Ignoring the intentions of the other agents in dynamic environments can lead to risky or over-conservative plans. In this work, we model the motion planning problem as a partially observable Markov decision process (POMDP) and propose an online system that combines an intent recognition algorithm and a POMDP solver to generate risk-bounded plans for the ego vehicle navigating with a number of dynamic agent vehicles. The intent recognition algorithm predicts the probabilistic hybrid motion states of each agent vehicle over a finite horizon using Bayesian filtering and a library of pre-learned maneuver motion models. We update the POMDP model with the intent recognition results in real time and solve it using a heuristic search algorithm which produces policies with upper-bound guarantees on the probability of near colliding with other dynamic agents. We demonstrate that our system is able to generate better motion plans in terms of efficiency and safety in a number of challenging environments including unprotected intersection left turns and lane changes as compared to the baseline methods.

Introduction

Driving in dynamic scenes, such as intersections and busy streets, is stressful because drivers need to operate their vehicles safely among other vehicles with uncertain motions. A study from (Choi 2010) shows that 40% of crashes that occurred in the United States in 2008 were intersection-related crashes. Among these intersection crashes, false assumption of other’s action and misjudgment of the gap or other’s speed are responsible for 8.4% and 5.5% of the critical reasons, respectively. These statistics indicate that driving can be safer and less stressful if the intentions of the other vehicles can be recognized accurately.

Since the first DARPA grand challenge in 2004, there have been many advances in the development of motion planning techniques (Paden et al. 2016) that can be used in intelligent vehicle systems to either assist the driver or take

*This work was partially supported by the Toyota Research Institute (TRI). However, this article solely reflects the opinions and conclusions of its authors and not TRI or any other Toyota entity. Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Two motivating scenarios, where the ego vehicle in yellow aims to go forward without colliding with the agent vehicle in red. Motion planners that ignore the intentions of the agent vehicles or assume fixed a intention distribution will produce results that are either risky or conservative.

over the control completely. Despite their success, many motion planners work only with static obstacles or assume simple intentions for the dynamic obstacles, such as moving at a constant velocity towards a fixed direction, which can lead to risky or conservative plans. In a recent work (Huang et al. 2018), the authors add uncertainty to the intentions of the moving obstacles such as nearby agent vehicles by assigning a probability distribution over what maneuver each agent vehicle will take, and propose a method that is able to produce risk-bounded motion plans for the ego vehicle that upper bound the probability of near collision with other agent vehicles taking stochastic actions. The work is limited due to its assumption on a fixed distribution over the agent vehicle intentions.

To motivate this limitation, we show two examples that can be seen in every driving experience in Fig. 1. In the first case, the autonomous ego vehicle is driving in the left lane behind a slow agent vehicle driving in the right lane. It is a fair assumption that the agent vehicle is more likely going to keep driving in the same lane, so the planner can generate a collision-free plan asking the ego vehicle to go forward and pass the slow vehicle. However, the same assumption fails in the second scenario, where there are two agent vehicles in the right lane, so there is a high likelihood that the following agent vehicle will make a lane change to the left especially when it is faster than the leading vehicle, and the

plan of passing it will lead to crash. On the other hand, if we assume that the agent vehicle is always more likely going to make a lane change to the left, the planner will slow down the ego vehicle in the second scenario in order to avoid the crash, but result in the ego vehicle acting too conservatively in the first scenario and failing to make the pass. Therefore, it is critical to dynamically infer the future motion of each surrounding agent vehicle in real-time, which requires a fast intent recognition ability.

We consider the vehicle motion as a sequence of movements depending on the high-level intended maneuver by the driver. Therefore, our proposed intent recognition algorithm is divided into two steps, where we first infer the driver maneuver and then predict a sequence of hybrid vehicle states including discrete maneuver and continuous trajectory over a finite horizon. The discrete maneuver prediction allows the planner to update the probability distribution over the agent vehicle maneuvers, and the continuous trajectory prediction helps to compute the near collision risk and identify risky plans. As studied in (Lefèvre, Vasquez, and Laugier 2014), maneuver-based motion prediction approach produces reliable longer-term motion predictions with reasonable computational time as compared to physics-based approaches and interaction-aware approaches.

Due to the uncertainties in the motions of human-driven vehicles, we pre-learn a compact motion representation called Probabilistic Flow Tube (PFT) (Dong and Williams 2012) from demonstrating trajectories to capture human-like driver styles and uncertainties for each maneuver. A library of pre-learned PFTs can be used to estimate the current maneuver as well as predict the probabilistic motion of each agent vehicle.

Given the probabilistic prediction results, we can update the belief state dynamically in a POMDP framework and solve it with an existing solver. Compared to many other intention-aware POMDP planners, our method is able to generate risk-bounded plans that provide guarantees on the risk defined as the probability of near collisions with other dynamic agents by augmenting the POMDP model with chance constraints and solving it with a heuristic search algorithm similar to (Santana, Thiébaux, and Williams 2016).

Our contributions are three-fold. First, we present a fast and accurate system for recognizing agent vehicle's intention, in terms of its maneuver and trajectory, using Bayesian filtering and a set of pre-learned maneuver motion models that consider maneuver uncertainties and encode different driver styles. Second, we present an intention-aware conditional motion planner that produces risk-bounded plans in dynamic driving scenes by considering the stochastic motions of other dynamic agents and the upper bound on the risk constraints. Lastly, we show that our planner can produce policies that echo different driver styles in an uncertain environment by adjusting the upper bound probability for the risk constraints in the POMDP model.

Related Work

In this section, we discuss state-of-the-art works relevant to our work in three areas: motion planning, motion prediction, and motion representation.

Motion Planning with Moving Obstacles

Vehicle motion planning problems with moving obstacles have been studied for decades. Many approaches use search-based methods to find a feasible path to the goal location by considering future motions of moving obstacles (Rudenko, Palmieri, and Arras 2017; Bansal et al. 2018; Ajanovic et al. 2018). However, these methods rely on simple prediction models, such as vehicle driving at current velocity and staying in the same lane (Ajanovic et al. 2018), and goal states of the obstacles are known a priori (Rudenko, Palmieri, and Arras 2017; Bansal et al. 2018). Such models are not realistic in planning for long-horizon where the obstacles can change both high-level goals and low-level paths over time.

A widely used framework for motion planning with dynamic obstacles is called Partially Observable Markov Decision Process (POMDP), which provides a systematic model that incorporates the uncertainty in the environment including the sensor noise and stochastic future motions of surrounding obstacles. Approaches in (Bandyopadhyay et al. 2013; Ulbrich and Maurer 2013; Brechtel, Gindele, and Dillmann 2014; Bai et al. 2015; Luo et al. 2018) model the intention-aware motion planning problem with POMDP where the intentions of the dynamic obstacles, such as vehicles and pedestrians, are represented by a belief state, and solve for the POMDP instance using approximation methods due to the complexity of the search space. Instead of solving the complete POMDP instance directly, (Chen et al. 2016) decouples the problem into multiple MDP instances by assuming fixed or deterministic intention. Additionally, Bayesian reinforcement learning (Wang et al. 2012; Hoang and Low 2013) and deep reinforcement learning (Everett, Chen, and How 2018) are proposed to enable safe and efficient motion planning in dynamic environments. These methods are limited in either working in discrete space or requiring a large amount of training data and computational resources.

Despite the success of POMDP-based approaches, they fail to provide a plan with a guarantee over the probability of success, such as avoiding near collision with moving agents in a dynamic environment. An extension called constrained POMDP is introduced to model risk explicitly and generate plans with bounded-risks as the world changes (Undurti and How 2010; Poupart et al. 2015). Using a similar idea, a chance-constrained POMDP (CC-POMDP) is proposed to include a more flexible definition of risk allocation and studied in different problem domains (Santana, Thiébaux, and Williams 2016). Unfortunately, the model assumes fixed or deterministic intentions of the dynamic agents, which can fail in many realistic driving scenarios. Therefore, we decide to improve the chance-constrained model by updating the belief state dynamically with a real-time intent recognition system, and solve it by a heuristic search algorithm named Risk-Bounded AO* (RAO*) (Santana, Thiébaux, and Williams 2016).

Maneuver-based Motion Prediction

Maneuver-based prediction algorithms predict the future trajectory of a target vehicle by estimating and predicting a se-

quence of high-level maneuver actions.

Vehicle maneuvers can be estimated and predicted using many approaches including threshold-based classifiers (Houenou et al. 2013), Bayesian Networks (Schreier, Willert, and Adamy 2014), Hidden Markov Model (Li et al. 2015), logistic regression (Klingelschmitt et al. 2014), Gaussian process regression (Tran and Firl 2013), and deep neural networks (Deo and Trivedi 2018). In this work, we focus on methods that output probability distributions over both discrete maneuvers and continuous trajectories.

The maneuver estimation results can help predict continuous trajectories over longer horizons. For example, (Houenou et al. 2013) generates a set of long-term trajectories based on the detected maneuver and selects the one with the minimum cost. In (Deo and Trivedi 2018), the authors train a deep neural network to generate a probability distribution over future locations conditioned on the states of the vehicle and the estimated maneuver. In (Schreier, Willert, and Adamy 2014), a hand-crafted motion model is used to predict future trajectories given the estimated maneuver. Similarly, our method uses a pre-learned motion model to generate probabilistic predictions on the target vehicle’s future positions.

Vehicle Motion Representation

Different motion representations have been used in the vehicle motion prediction domain. Examples include poly-lines (Houenou et al. 2013), maneuver-specific models (Schreier, Willert, and Adamy 2014), Gaussian mixture models (Deo and Trivedi 2018), and funnels (Tobenkin, Manchester, and Tedrake 2011), etc. Since some trajectories are more likely to occur than others when a human driver executes a particular maneuver, we favor probabilistic motion representations, which are useful to incorporate the uncertainties from human driver behavior and to compute the near collision probability in a chance-constrained framework.

The work in (Dong and Williams 2012) proposes a representation called Probabilistic Flow Tubes that models human-like actions in humanoid robotics tasks by computing common characteristics defined over a time interval from a set of demonstrating examples.

We decide to use PFT to model probabilistic vehicle motions because of a few reasons. First, it is a compact representation to capture both temporal and spatial aspects of the maneuver motions, which is parameterized by a sequence of Gaussian distributions defined by means and covariances over different time steps. Second, it learns human behavior through a number of demonstrating examples. In the vehicle domain, the PFT can capture different driving styles efficiently given enough demonstrating trips. Although many approaches estimate the driving styles more explicitly from onboard vehicle sensors (Martinez et al. 2018), PFT is able to incorporate driver styles inherently in its representation.

Problem Formulation

Our goal is to generate a motion plan for a controllable ego vehicle that navigates in a dynamic environment with other uncontrollable agent vehicles with stochastic motions.




Figure 2: Architecture diagram for our risk-bounded intention-aware motion planning system.

The input includes a library of pre-learned maneuver motion models and a sequence of observed vehicle states. The output is a policy that achieves the minimum cost while maintaining an upper bound on risk defined as the probability of near colliding with other agent vehicles. The overall architecture diagram of our proposed planning system is depicted in Fig. 2.

We start by modeling the vehicle motion planning problem with a *CC-POMDP* instance (Santana, Thiébaux, and Williams 2016) that defines the problem of planning with multiple agents including a controllable agent \mathcal{R}_0 and a number of n uncontrollable agents $\mathcal{R}_1, \dots, \mathcal{R}_n$ with no communication in between. The vehicle-specific instance is a tuple $\mathcal{P} = \langle \mathcal{S}, \mathcal{A}, \mathcal{O}, T, O, C, \mathcal{B}_0, H, \Delta \rangle$, where

- \mathcal{S} is a set of hybrid states for each vehicle including the positions and intended maneuvers, where the ego vehicle has no intended maneuvers.
- \mathcal{A} is the set of maneuvers that can be taken by the ego vehicle.
- \mathcal{O} is a set of observations on the vehicle positions.
- $T : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is a state transition function between hybrid states. The transition function depends not only on the action chosen by \mathcal{R}_0 , but also on the intention of each uncontrollable vehicle.
- $O : \mathcal{S} \times \mathcal{O} \rightarrow \mathbb{R}$ is an observation function that models the probability of observing a sequence of agent vehicle positions given the state information such as intended maneuvers.
- $C : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is a cost function associated with each action in \mathcal{A} .
- \mathcal{B}_0 is the initial belief state over hybrid states.
- H is the finite planning horizon.
- Δ is an upper bound probability for the risk constraint, where the risk is defined as the probability of \mathcal{R}_0 near colliding with any of the agent vehicles up to the planning horizon.

Given the instance, the objective is to find an optimal, deterministic, and chance-constrained policy π^* at time step k , such that

$$\pi^* = \arg \min_{\pi} \mathbb{E} \left[\sum_{t=k}^{k+H} C(s_t, a_t) \middle| \pi \right], \quad (1)$$

subject to

$$er(\mathcal{B}_k | \pi^*) = Pr \left(\bigvee_{i=k}^{k+H} \text{Collision}_i | \mathcal{B}_k, \pi^* \right) \leq \Delta, \quad (2)$$

where er is the execution risk of a policy π measured from the current belief state \mathcal{B}_k over planning horizon H .

We assume the intended maneuvers of the agent vehicles to be partially observable, and introduce a Probabilistic Hybrid Automaton (PHA) including as a 6-tuple $\mathcal{H} = \langle \mathbf{x}, \mathbf{w}, F, T_d, \mathcal{X}_d, T_s \rangle$ (Hofbaur and Williams 2002) to estimate the hybrid states for each uncontrollable agent vehicle:

- $\mathbf{x} = \mathbf{x}_d \cup \mathbf{x}_c$ are the state variables, where $\mathbf{x}_d \in \mathcal{X}_d$ includes an intended maneuver type and a maneuver clock, and \mathbf{x}_c denotes continuous state variables as global positions (x, y) of the target vehicle.
- $\mathbf{w} = \mathbf{u}_d \cup \mathbf{o} \cup \mathbf{v}_c$ are the interface variables, where $\mathbf{u}_d = \emptyset$ denotes the set of discrete input variables because we assume no access to the control over the uncontrollable agents, $\mathbf{o} = \{z_x, z_y\}$ includes the observed vehicle locations, and $\mathbf{v}_c = \emptyset$ assumes the vehicle positions are fully observable.
- $F = \mathcal{F}_c \cup \mathcal{G}_c$ specifies the continuous evolution of the automaton. For each discrete state $x_d \in \mathcal{X}_d$, F_c specifies the discrete-time dynamic evolution of the vehicle, and is represented by a Probabilistic Flow Tube (PFT) with sampling period T_s . More details of PFT will be explained in Section APPROACH. In addition, $\mathcal{G}_c(a) = a$ is a perfect observation function for any continuous variable a .
- T_d specifies a set of transition functions for each discrete state. Each transition function has an associated guard condition determined by the continuous state variables.

After defining a PHA specifically for each uncontrollable vehicle, we introduce a sub-problem of intent recognition as generating the probability distribution over the hybrid state over a finite predicting horizon h that consists of continuous trajectory \mathbf{x}_c together with the discrete state \mathbf{x}_d given a sequence of observed history vehicle locations $\mathbf{o}_{k-w+1:k}$ with length w :

$$P(\mathbf{x}_{c,k:k+h}, \mathbf{x}_{d,k:k+h} | \mathbf{o}_{k-w+1:k}). \quad (3)$$

Note that the predicting horizon h is different from the planning horizon H defined in the POMDP model. The predicting horizon refers to the number of time steps ahead of the current time step for intent recognition, and the planning horizon refers to the number of actions the planner needs to look ahead. Usually, the predicting horizon is the product of the planning horizon and the number of clock cycles contained in the maneuver action.

Finally, the intent recognition results can be used to update the belief state and the transition function T in the CC-POMDP instance, as explained in the next section.

Approach

This section introduces three major modules in our intention-aware motion planning system as illustrated in Fig. 2: a *Motion Model Generator* generating a library of Probabilistic Flow Tubes (PFTs) that model the continuous evolution of high-level vehicle maneuvers, an *Intent Recognizer* that predicts probabilistic maneuver states and continuous vehicle positions over a finite horizon for each uncontrollable agent vehicle, and a *Risk-bounded POMDP Planner* that considers the agent car intentions and produces the desired maneuver plan for the ego car.

Probabilistic Maneuver Motion Modeling

Probabilistic Flow Tubes (PFTs) are used to model the probabilistic continuous motion for autonomous systems. As described in (Dong and Williams 2012), a PFT is a compact representation for continuous trajectories with common characteristics. It is parameterized by a set of means and covariances that correspond to each time step of the common trajectories.

Since the trajectories are generated from human demonstrations, PFTs are useful to model realistic human behaviors and incorporate different driver styles. The algorithm to generate a PFT for each discrete maneuver $M \in \mathcal{M}$ given a set of demonstrating trajectories \mathcal{D}_M is shown in Alg. 1. Since the trajectories associated with the same maneuver may have different lengths, we first perform dynamic time warping (DTW) on the input trajectories to temporally align them (Myers, Rabiner, and Rosenberg 1980). Assuming all aligned trajectories have a length of L , we compute the means and covariances among all aligned trajectories at time steps from 1 to L .

Algorithm 1 Generate a Probabilistic Flow Tube for M

Input $M, \mathcal{D}_M = \{D_M^{(n)}\}_{n=1 \dots N}$, a set of N trajectories

Output L_M

- 1: **procedure** GENERATEPFT(\mathcal{D}_M)
 - 2: $\mathcal{D}'_M = \text{DTW}(\mathcal{D}_M)$
 - 3: **for** $i = 1$ to L **do**
 - 4: $\mathbf{x}_i \leftarrow (D_M^{(1)}[i], \dots, D_M^{(N)}[i])$
 - 5: $\mu[i] = \text{mean}(\mathbf{x}_i)$
 - 6: $\Sigma[i] = \text{covariance}(\mathbf{x}_i)$
 - 7: $L_M = (\mu, \Sigma)$
 - 8: **return** L_M
-

Therefore, given a discrete state including a maneuver label M and a maneuver clock i , the probability of the continuous state is given as a Gaussian distribution:

$$p(\mathbf{x}_c | \mathbf{x}_d = (M, i)) = \mathcal{N}(\mu_M[i], \Sigma_M[i]). \quad (4)$$

In addition, similar to (Dong and Williams 2012), the likelihood of observing a continuous past trajectory $\mathbf{o}_{k-w+1:k}$ given a maneuver label M and a maneuver clock i is the product of probability densities of each Gaussian distribution in the flow tube evaluated at the aligned points in the

observed vehicle trajectory. Since the trajectory is not necessarily aligned with the flow tube, we shift the end of the trajectory so that it is aligned with $\mu_M[i]$ and name the shifted trajectory $\mathbf{o}'_{k-w+1:k}$. Therefore, the likelihood is computed as:

$$\begin{aligned} p(\mathbf{o}_{k-w+1:k}|\mathbf{x}_d) &= (M, i) \\ &= \prod_{j=k-w+1}^k \mathcal{N}(\mathbf{o}'_j|\mu_M[i-k+j], \Sigma_M[i-k+j]). \end{aligned} \quad (5)$$

Probabilistic Hybrid State Prediction

We divide the probabilistic hybrid predictor into two steps. In the first step, we estimate the discrete maneuver state at the current time step k using Bayesian filtering and a library of pre-learned PFTs. Next, given the discrete estimation result, we compute the probability distribution over the future hybrid states using a recursive approach.

Discrete State Estimation Estimating the hybrid state in a PHA is considered as a *hybrid mode estimation problem* (Hofbaur and Williams 2002). Since the continuous vehicle locations are assumed to be fully observable, our goal is simplified to generate the probability distribution over the discrete maneuver state \mathbf{x}_d at the current time k given a PHA and an observation sequence $\mathbf{o}_{k-w+1:k}$ (simplified as \mathbf{o}) with length w .

The probability distribution can be computed by multiplying the observation probability using Equation 5 and the prior over the discrete states, where the prior is initialized to be a uniform distribution:

$$p(\mathbf{x}_{d,k}|\mathbf{o}) \propto p(\mathbf{o}|\mathbf{x}_{d,k})p(\mathbf{x}_{d,k}). \quad (6)$$

Hybrid State Prediction Given the estimation on the discrete state at current step k , we wish to predict the future hybrid vehicle states over a predicting horizon of h . To start, we compute the probability distribution over the next hybrid step at $k+1$:

$$p(\mathbf{x}_{d,k+1}, \mathbf{x}_{c,k+1}|\mathbf{o}) = p(\mathbf{x}_{d,k+1}|\mathbf{o})p(\mathbf{x}_{c,k+1}|\mathbf{x}_{d,k+1}, \mathbf{o}) \quad (7)$$

The first term $p(\mathbf{x}_{d,k+1}|\mathbf{o})$ in Equation 7 can be computed by Equation 8, where $p(\mathbf{x}_{d,k+1}|\mathbf{x}_{d,k}, \mathbf{o})$ can be computed using the discrete transition function T_d and its associated guard condition determined by observations \mathbf{o} , as specified in the PHA model, and $p(\mathbf{x}_{d,k}|\mathbf{o})$ is given by Equation 6:

$$p(\mathbf{x}_{d,k+1}|\mathbf{o}) = \sum_{\mathbf{x}_{d,k}} p(\mathbf{x}_{d,k+1}|\mathbf{x}_{d,k}, \mathbf{o})p(\mathbf{x}_{d,k}|\mathbf{o}). \quad (8)$$

The second term $p(\mathbf{x}_{c,k+1}|\mathbf{x}_{d,k+1}, \mathbf{o})$ in Equation 7 represents the probability distributions of vehicle locations given the predicted discrete maneuver state, which can be computed by looking up the PFT associated with the discrete maneuver state using Equation 4.

Applying Equations 7 and 8 recursively, we can estimate the future maneuver sequences as well as continuous states over a predicting horizon of h for the target vehicle.

The size of possible values for $\mathbf{x}_{d,k:k+h}$ increases exponentially with h , which causes the POMDP solver to

search over a very large search space. To reduce the size of possible discrete states, we set the probability of very low likely events (e.g., with probability smaller than a predefined threshold ϵ) at each predicting step to be 0.

Intention-Aware Risk-Bounded Motion Planning

In this section, we present an online intention-aware planning algorithm by combining a heuristic forward search method with intent recognition results, as outlined in Algorithm 2.

Algorithm 2 Online Intention-Aware Motion Planning

Input CC-POMDP \mathcal{P} , PHA \mathcal{H} , library of PFTs \mathcal{L} .

Output None.

```

1:  $\mathcal{P}_0 \leftarrow \mathcal{P}$ 
2:  $\mathcal{B}_0 \leftarrow \text{GetInitialBeliefState}(\mathcal{P}_0)$ 
3:  $k \leftarrow 1$ 
4: while True do
5:    $\mathbf{o}_k \leftarrow \text{ObserveCurrentVehicleLocations}()$ 
6:    $\mathbf{o}_{k-w+1:k} \leftarrow \text{UpdateObservation}(\mathbf{o}_k, \mathbf{o}_{k-w:k-1})$ 
7:    $\mathbf{p}(\mathbf{x}_{d,k}) \leftarrow \text{DiscreteEstimation}(\mathbf{o}_{k-w+1:k}, \mathcal{H}, \mathcal{L}, \mathcal{B}_{k-1})$ 
8:    $\mathcal{B}_k \leftarrow \text{UpdateBeliefState}(\mathbf{o}_k, \mathbf{p}(\mathbf{x}_{d,k}))$ 
9:   if AtGoal( $\mathcal{B}_k$ ) then
10:    Output success.
11:    break
12:    $\mathbf{p}(\mathbf{x}_{d \rightarrow}, \mathbf{x}_{c \rightarrow}) \leftarrow \text{Predict}(\mathbf{o}_{k-w+1:k}, \mathcal{H}, \mathcal{L}, \mathbf{p}(\mathbf{x}_{d,k}))$ 
13:    $T_k \leftarrow \text{UpdateTransition}(\mathbf{p}(\mathbf{x}_{d \rightarrow}, \mathbf{x}_{c \rightarrow}))$ 
14:    $\mathcal{P}_k \leftarrow \text{UpdateModel}(\mathcal{P}_{k-1}, \mathcal{B}_k, T_k)$ 
15:    $\pi_k \leftarrow \text{RAO}^*(\mathcal{P}_k)$ 
16:   Execute(GetAction( $\pi_k, \mathcal{B}_k$ ))
17:    $k \leftarrow k + 1$ 

```

We first initialize the necessary variables given the inputs to the problem (Line 1-3). At each execution step k , we begin by updating the observation sequence by collecting the latest vehicle positions through perception sensors (Line 5-6). We assume that we always have enough observations from the past. Since the major uncertainty in this work comes from the intentions of agent vehicles, we assume that the perception sensors are perfect and thus the vehicle states are fully observable.

Next, we estimate the possible maneuvers for each agent vehicle by taking the prior from the current belief state and computing the posterior using Equation 6 (Line 7). The estimation results and the most recent observed vehicle positions are used to update the current belief state (Line 8). If the belief state satisfies the goal condition, such as the ego vehicle arriving at a predefined location or region, we output success and end the planning process. Otherwise, we need to continue to predict the future hybrid states $\mathbf{p}(\mathbf{x}_{d,k:k+h}, \mathbf{x}_{c,k:k+h})$ (abbreviated as $\mathbf{p}(\mathbf{x}_{d \rightarrow}, \mathbf{x}_{c \rightarrow})$) for all agent vehicles (Line 12). The predictions are incorporated into the transition function, which keeps information on how the agent vehicles will move in the future (Line 13-14).

Once the POMDP is updated with the latest intent recognition results, our system can utilize any off-the-shelf POMDP solver to find a feasible policy. Because of the bounding risk requirement in our problem, we use a method

called risk-bounded AO* (RAO*) (Santana, Thiébaux, and Williams 2016; Huang et al. 2018) that finds the solution through heuristic forward search for CC-POMDP (Line 15), where the value heuristic is defined as the Euclidean distance to the goal location. To search for a risk-bounded policy, RAO* uses a second risk heuristic, based on execution risk, to prune the search space by removing the over-risky branches that violate the risk constraints. The execution risk, as defined in Equation 2, is computed as the probability of near collisions between our ego vehicle and agent vehicles, based on the intent recognition results.

In order to find solutions in real-time, progress has been made to speed up RAO* including adding preconditions for each action to further prune the search space and considering only the interaction with nearby vehicles to reduce the branching factor. Another possible way to improve computational time is to use an iterative RAO* algorithm that reuses the search graph during replanning (Huang et al. 2018). Unfortunately, the changing intentions of the agent cars restrict the amount of search graph that can be reused, which means in many cases we need to build a new graph from scratch. Therefore, for the sake of simplicity, RAO* is used and we show in the next section that it is sufficient to solve the vehicle motion planning problem with short horizons.

Given the output policy from the solver, our system can quickly look up the optimal action to execute (Line 16). The system would continue to iterate through the same steps in a loop until the goal condition is satisfied.

In general, the online intention-aware planning ability allows our system to make adjustments quickly and always have a risk-bounded plan when the surrounding agent vehicles change their intentions. This is a major advantage of our system as compared to other works using offline policies or assuming a fixed distribution over the vehicle intentions.

Experiments

In order to evaluate our system, we create an intersection unprotected left turn scenario and a dynamic lane change scenario in a state-of-the-art urban driving simulator called CARLA (Dosovitskiy et al. 2017). The simulator is able to produce a realistic driving experience in a small town and allows the users to create multiple vehicles based on their needs. Each vehicle can be controlled individually through steering wheel angle and throttle inputs collected from either a user input device such as a keyboard or a programmed software that generates the control given the environment information. In each scenario, we aim to test the different aspects of our system, as detailed in the rest of this section.

Intersection Left Turn

According to (Choi 2010), intersections are responsible for 40% of crashes occurred in the United States. Among all intersection crashes, unprotected left turns are one of the most common and dangerous ones. Therefore, we wish to verify our intention-aware planning system in a simplified scenario, where the ego car needs to make a left turn at an unprotected intersection when an agent car is approaching from the opposite direction. We aim to show that our system



(a) View of the ego car. (b) View of the agent car.

Figure 3: Intersection left turn scenario setup.

is able to handle this dynamic environment and outperform other baseline methods. A follow-up experiment demonstrates that our system is able to simulate different driving styles by tuning the upper bound on the risk-constraint.

Scenario Setup As shown in Fig. 3, we start the test by placing the yellow ego vehicle in front of an intersection, and the red agent vehicle on the opposite side. The ego car has a goal to turn left. At each step, it can choose to stop and yield to the agent car or perform a left turn. In order to encourage turning sooner than later, we add a higher cost to the stop action. On the other hand, the agent car chooses randomly from two maneuvers including going forward or slowing down to yield to the ego car. We pre-learn a library of PFTs representing the uncertainties of these two maneuvers by asking several human drivers to teleop the agent car for 100 times in the simulator using a Logitech steering wheel and pedal device. Furthermore, to maintain a reasonable computational time, we choose to use 0.0001 as the value of ϵ that filters out approximately 90% of possible discrete states. In addition, as described in the definition of PHA, we assume that the ego vehicle has access to the exact locations of all agent vehicles.

Results We start by investigating the performance of our intent-recognition algorithm by separately measuring the accuracy of maneuver estimation and trajectory prediction. To measure the maneuver estimation accuracy, we compute the percentage of correctly identified maneuvers among the testing set. Since the intent recognition algorithm outputs a probability distribution over possible maneuvers, we choose the most likely maneuver and compare it with the ground truth maneuver. To measure the motion prediction performance, we sample from the predicted trajectory distribution and compute the average end displacement error over 4.8 seconds among all testing samples. The predicting horizon is chosen as the time taken by the ego car to make a left turn. Among all the testing examples, our intent estimation has an average accuracy of 96.19%, and the average end displacement error is 2.02 meters. We observe that most maneuver misclassifications are due to some forward testing trajectories that contain a sequence of consecutive positions connected with small gaps. Such sequence confuses the maneuver estimator to produce a distribution that assigns more probability to the slow down maneuver than to the forward maneuver, as well as causes a large trajectory prediction error.

Next, we want to check the performance of our proposed

	Success Rate [%]	Completion Time [s]
Conservative	100.00	9.60
Risky	56.80	4.80
AccelerationBased	82.10	8.87
Ours ($\Delta = 0.1\%$)	100.00	8.64
Conservative	100.00	26.74
Risky	28.60	19.20
Ours ($\Delta = 0.1\%$)	100.00	24.67
Ours ($\Delta = 1\%$)	99.80	23.19

Table 1: Performance of different intention-aware planners. Top half: unprotected left turn. Bottom half: lane change.

intention-aware motion planning approach. Since this work focuses on extending a risk-bounded motion planner with intention recognition capabilities, we compare our approach with three baseline planners that have simple prediction models. The first baseline planner assumes that the agent car will choose each maneuver with equal likelihood. The second baseline assumes that the agent vehicle is slowing down with 100% probability, and thus always produces a plan to turn immediately. The first two baseline planners all use the RAO* planner to find the policy offline. Instead of using a risk-bounded planner, the third baseline planner chooses the action for the ego car based on the acceleration of the agent car. If the acceleration becomes smaller than a negative threshold, it assumes the agent car is slowing down and executes the ego car to make a left turn.

We test each planner with 1000 trials. Two measurements are used to quantify the performance in terms of safety and efficiency: the percentage of collision-free turns and the average time to make turns among all collision-free turns.

Top four rows in Table 1 summarize the performance of our planner and three baseline planners in the unprotected left turn scenario. The first baseline method is quite successful in avoiding all potential near collisions by using a conservative assumption. However, they sacrifice the efficiency by keeping the ego car waiting even when the agent car slows down and yields. On the other hand, the second baseline is efficient in making turns with the smallest amount of completion time, but leads to crashes almost half of the time because of its over-risky assumption. The acceleration-based method works in most of the cases, but it is vulnerable in examples where the agent car decelerates slightly in the middle of a forward maneuver. Compared to the methods above, our model is able to find safe plans in all cases by allowing only a small risk bound of 0.1%, and completes the tasks more efficiently by detecting the slow-down intention of the agent car quickly and reliably.

We further analyze the performance of our planner by changing the risk bound Δ . As plotted in Fig. 4, when the upper risk bound is relatively small, the planner can guarantee to find collision-free policies for the ego car because it is very rare that our estimator classifies a forward maneuver into a slow down maneuver for the oncoming agent car




Figure 4: The performance of our planner in the intersection left turn scenario as a function of the upper risk bound Δ , where the x -axis is plotted in the log scale.

with great confidence. However, the smaller the risk bound, the more time it takes the ego car to complete the task as the planner needs more time to be certain that it is within the safe bound to make a left turn. When the risk bound becomes larger, the success rate starts to drop because in some cases the intent recognizer provides inaccurate estimations, which leads our planner that takes large risk bound to choose the wrong action for the ego vehicle. When the risk bound is set to one, the behavior of the planner is essentially identical to the second baseline planner that takes risky decisions.

Therefore, we can tune the *style* of our planner using different risk bound values, where a small risk bound leads to safer behaviors, and a large risk bound produces a risky planner that is willing to sacrifice safety in order to achieve better efficiency. The planner style can be either determined offline by a user, or adjusted online during a planning episode based on different driving conditions.

Lane Change with Multiple Agent Vehicles

In this scenario, we aim to test our planner in a more complicated scenario, where the ego vehicle interacts with more than one agent car on a busy road and needs to make lane changes to achieve better efficiency. In addition to comparing the performances with baseline methods, we analyze the computational time of our planner by varying the number of agent cars and the planning horizon.

Scenario Setup As shown in Fig. 5, we place the ego vehicle in the middle of a two-lane road at the beginning of each test. It can choose to change lane or go forward, and its goal is to reach the end of the road as soon as possible. In addition, we place three agent vehicles randomly next to the ego vehicle, each of which can choose to change lane or go forward. The agent cars are designed to move more slowly than the ego car, which encourages the ego car to bypass them by changing lanes. Again, a set of PFTs for all agent vehicle maneuvers are learned from human demonstrations, and we use the same ϵ value as in the previous scenario.

The starting location of each agent vehicle is chosen so



Figure 5: A lane change test example, where the yellow vehicle is the ego car, and the reds are the agent cars.

that they can create potential risks to the ego car. Fig. 5 shows one testing example, where an agent vehicle is placed behind the yellow ego car in the adjacent lane to prevent the ego car from making a lane change at the beginning of the test. The other two agent vehicles are placed in the front on different lanes to create potential risks. From many experiments, we find that three cars are sufficient to cover most dangerous lane change scenarios.

Results In this scenario, the intent recognition provides accurate results with an average maneuver estimation accuracy of 98.5%, and an average trajectory prediction error of 2.23 meters. A better maneuver estimation performance is achieved compared to the left turn scenario as lane change and forward maneuvers can be more easily distinguished in the lateral axis. However, the trajectory prediction error increases because a misclassified change-lane maneuver will result in a trajectory prediction in the opposite direction.

Similar to the intersection left turn scenario, we test our system and two baseline planners that have simple prediction models with 1000 trials. The results are summarized in the bottom half of Table 1. Our planner is able to produce more efficient plans than the conservative planner while maintaining safety all the time if the risk bound is set to be small. In addition, we can make an aggressive planner to further reduce the completion time by increasing the risk bound. Although the success rate drops, our planner is able to respect the risk bound with guarantees.

We choose a planning horizon of two because the number of possible future states for the agent vehicles increases exponentially with the planning horizon, and the planner needs to compute the near collision risk for each possible future state. In Fig. 6, we plot the planning time in log scale as a function of the planning horizon for different numbers of agent vehicles, where each vehicle can choose from three maneuvers and ϵ is fixed. The dashed black line indicates the maximum allowed planning time of 0.2 seconds if we want to keep the control frequency at 5Hz. With three agent vehicles and a planning horizon of two, it takes around 0.191 seconds for the planner to find a solution, which satisfies the planning time constraint. Although a planning horizon of two is not ideal, our ego vehicle is able to make adjustments quickly if there exists possible risk further ahead thanks to the online capability of our planner. Furthermore, we ob-




Figure 6: Planning time using RAO* planner as a function of the planning horizon H and the number of agent vehicles n . The planning time is shown in log scale.

serve that in many cases the ego vehicle only needs to interact with one or two vehicles. Therefore, we can adjust the planning horizon based on the number of interacting agents to improve the performance of our planner.

Conclusion

In conclusion, we propose an online risk-bounded planning system that considers the intentions of surrounding dynamic agents, including their intended maneuvers and future motions. The intentions are recognized by a Bayesian filter based on a probabilistic hybrid automaton that models the agent vehicle behaviors. We integrate the intent recognition results into a POMDP model and solve it with a heuristic forward search method with upper bound guarantees on the near collision probability with other agent vehicles. Our system is demonstrated to work in two challenging dynamic environments in real time and outperform other baseline methods. We also show a further analysis of how the upper bound on risk constraint can change the behavior of our planner. We believe that both risk guarantee and intent recognition capability are necessary for safe-critical motion planning tasks in a dynamic environment, and our work provides a promising framework in that direction. One possible future work is to test our work in real systems and more complicated scenarios. In addition, we would like to apply our method to other domains such as interactive robotics manipulation.

References

- Ajanovic, Z.; Lacevic, B.; Shyrokau, B.; Stolz, M.; and Horn, M. 2018. Search-based optimal motion planning for automated driving. *arXiv preprint arXiv:1803.04868*.
- Bai, H.; Cai, S.; Ye, N.; Hsu, D.; and Lee, W. S. 2015. Intention-aware online pomdp planning for autonomous driving in a crowd. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, 454–460. IEEE.
- Bandyopadhyay, T.; Won, K. S.; Frazzoli, E.; Hsu, D.; Lee, W. S.; and Rus, D. 2013. Intention-aware motion planning. In *Algorithmic foundations of robotics X*. Springer. 475–491.

- Bansal, S.; Cosgun, A.; Nakhaei, A.; and Fujimura, K. 2018. Collaborative planning for mixed-autonomy lane merging. *arXiv preprint arXiv:1808.02550*.
- Brethel, S.; Gindele, T.; and Dillmann, R. 2014. Probabilistic decision-making under uncertainty for autonomous driving using continuous pomdps. In *Intelligent Transportation Systems (ITSC), 2014 IEEE 17th International Conference on*, 392–399. IEEE.
- Chen, M.; Frazzoli, E.; Hsu, D.; and Lee, W. S. 2016. Pomdp-lite for robust robot planning under uncertainty. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on*, 5427–5433. IEEE.
- Choi, E. 2010. Crash factors in intersection-related crashes: An on-scene perspective (no. dot hs 811 366). *US DOT National Highway Traffic Safety Administration*.
- Deo, N., and Trivedi, M. M. 2018. Multi-modal trajectory prediction of surrounding vehicles with maneuver based lstms. *arXiv preprint arXiv:1805.05499*.
- Dong, S., and Williams, B. 2012. Learning and recognition of hybrid manipulation motions in variable environments using probabilistic flow tubes. *International Journal of Social Robotics* 4(4):357–368.
- Dosovitskiy, A.; Ros, G.; Codevilla, F.; Lopez, A.; and Koltun, V. 2017. CARLA: An open urban driving simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*, 1–16.
- Everett, M.; Chen, Y. F.; and How, J. P. 2018. Motion planning among dynamic, decision-making agents with deep reinforcement learning. *arXiv preprint arXiv:1805.01956*.
- Hoang, T. N., and Low, K. H. 2013. A general framework for interacting bayes-optimally with self-interested agents using arbitrary parametric model and model prior. In *IJCAI*, 1394–1400.
- Hofbauer, M. W., and Williams, B. C. 2002. Mode estimation of probabilistic hybrid systems. In *International Workshop on Hybrid Systems: Computation and Control*, 253–266. Springer.
- Houenou, A.; Bonnifait, P.; Cherfaoui, V.; and Yao, W. 2013. Vehicle trajectory prediction based on motion model and maneuver recognition. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, 4363–4369. IEEE.
- Huang, X.; Jasour, A.; Deyo, M.; Hofmann, A.; and Williams, B. C. 2018. Hybrid risk-aware conditional planning with applications in autonomous vehicles. In *Decision and Control (CDC), 2018 IEEE International Conference on*. IEEE.
- Klingelschmitt, S.; Platho, M.; Groß, H.-M.; Willert, V.; and Eggert, J. 2014. Combining behavior and situation information for reliably estimating multiple intentions. In *Intelligent Vehicles Symposium (IV), 2014 IEEE*, 388–393. IEEE.
- Lefèvre, S.; Vasquez, D.; and Laugier, C. 2014. A survey on motion prediction and risk assessment for intelligent vehicles. *Robomech Journal* 1(1):1.
- Li, G.; Li, S. E.; Liao, Y.; Wang, W.; Cheng, B.; and Chen, F. 2015. Lane change maneuver recognition via vehicle state and driver operation signals—results from naturalistic driving data. In *Intelligent Vehicles Symposium (IV), 2015 IEEE*, 865–870. IEEE.
- Luo, Y.; Cai, P.; Bera, A.; Hsu, D.; Lee, W. S.; and Manocha, D. 2018. Porca: Modeling and planning for autonomous driving among many pedestrians. *IEEE Robotics and Automation Letters* 3(4):3418–3425.
- Martinez, C. M.; Heucke, M.; Wang, F.-Y.; Gao, B.; and Cao, D. 2018. Driving style recognition for intelligent vehicle control and advanced driver assistance: A survey. *IEEE Transactions on Intelligent Transportation Systems* 19(3):666–676.
- Myers, C.; Rabiner, L.; and Rosenberg, A. 1980. Performance tradeoffs in dynamic time warping algorithms for isolated word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 28(6):623–635.
- Paden, B.; Čáp, M.; Yong, S. Z.; Yershov, D.; and Frazzoli, E. 2016. A survey of motion planning and control techniques for self-driving urban vehicles. *IEEE Transactions on Intelligent Vehicles* 1(1):33–55.
- Poupart, P.; Malhotra, A.; Pei, P.; Kim, K.-E.; Goh, B.; and Bowling, M. 2015. Approximate linear programming for constrained partially observable markov decision processes. In *AAAI*, 3342–3348.
- Rudenko, A.; Palmieri, L.; and Arras, K. O. 2017. Predictive planning for a mobile robot in human environments. In *IEEE International Conference on Robotics and Automation (ICRA), Workshop on PlanRob*.
- Santana, P.; Thiébaux, S.; and Williams, B. 2016. Rao*: An algorithm for chance-constrained pomdps. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI'16)*, 3308–3314.
- Schreier, M.; Willert, V.; and Adamy, J. 2014. Bayesian, maneuver-based, long-term trajectory prediction and criticality assessment for driver assistance systems. In *Intelligent Transportation Systems (ITSC), 2014 IEEE 17th International Conference on*, 334–341. IEEE.
- Tobenkin, M. M.; Manchester, I. R.; and Tedrake, R. 2011. Invariant funnels around trajectories using sum-of-squares programming. *IFAC Proceedings Volumes* 44(1):9218–9223.
- Tran, Q., and Firl, J. 2013. Modelling of traffic situations at urban intersections with probabilistic non-parametric regression. In *Intelligent Vehicles Symposium (IV), 2013 IEEE*, 334–339. IEEE.
- Ulbrich, S., and Maurer, M. 2013. Probabilistic online pomdp decision making for lane changes in fully automated driving. In *Intelligent Transportation Systems (ITSC), 2013 IEEE 16th International Conference on*, 2063–2067. IEEE.
- Undurti, A., and How, J. P. 2010. An online algorithm for constrained pomdps. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, 3966–3973. IEEE.
- Wang, Y.; Won, K. S.; Hsu, D.; and Lee, W. S. 2012. Monte carlo bayesian reinforcement learning. *arXiv preprint arXiv:1206.6449*.