Improved Occlusion Scenario Coverage with a POMDP-based Behavior Planner for Autonomous Urban Driving

Chi Zhang[1], Florian Steinhauser[2], Gereon Hinz[3], Alois Knoll[4]

***Abstract*— Safely driving through various occlusion scenarios in urban environments, such as bus stops or crosswalks, is challenging for autonomous vehicles (AVs). Improving the ability to handle more occlusion scenarios in urban environments is paramount when using AVs as shuttle buses. An AV could experience deadlock situations in very heavy occlusion scenarios with the worst-case assumption that potential occluded road users could suddenly emerge using maximal allowed velocity. In this study, we address this issue with a partially observable Markov decision process (POMDP)-based behavior planner to improve the occlusion scenario coverage. We extend a phantom vehicle concept to include pedestrians to represent potential road users in risky occlusion areas. The appearance probability of phantom objects along with their future movement is inferred using map information and road topology. Finally, context-aware phantom road users are incorporated within a POMDP formulation, which is solved online by constructing a Monte Carlo tree with reachable state analysis. Various evaluation results indicate that the ego vehicle shows comfortable driving behavior, aiming to avoid unnecessary braking and acceleration when driving through challenging occlusion scenarios in urban areas, including crosswalks, bus stops, and intersections. Moreover, it does not lead to deadlock situations in heavily occluded scenarios.**

I. INTRODUCTION

For autonomous vehicles (AVs), a decision-making module plays a decisive role in enabling safe and socially acceptable driving in all traffic scenarios. This role is analogous to the brain of a human driver. The challenges of safe decision making include observing traffic rules and reasonably interacting with other traffic participants. Further, the uncertainties of perceived traffic information, such as noise in sensor measurements, uncertain motion predictions, and occluded objects in complex urban areas, need to be considered. Although some state-of-the-art AVs work well under sound conditions, handling these problems remains nontrivial.

Recent works have addressed the aforementioned challenges in decision making, focusing on occlusion due to static and dynamic obstacles in intersections. However, existing online behavior planning algorithms do not perform well in urban traffic scenarios where pedestrians are present, such as serving bus stops or approaching crosswalks (see Fig. 1b). These scenarios are challenging because a decision-making system must consider traffic participants in the observable region and reason about the potential road users and their movement in invisible areas based on context information.

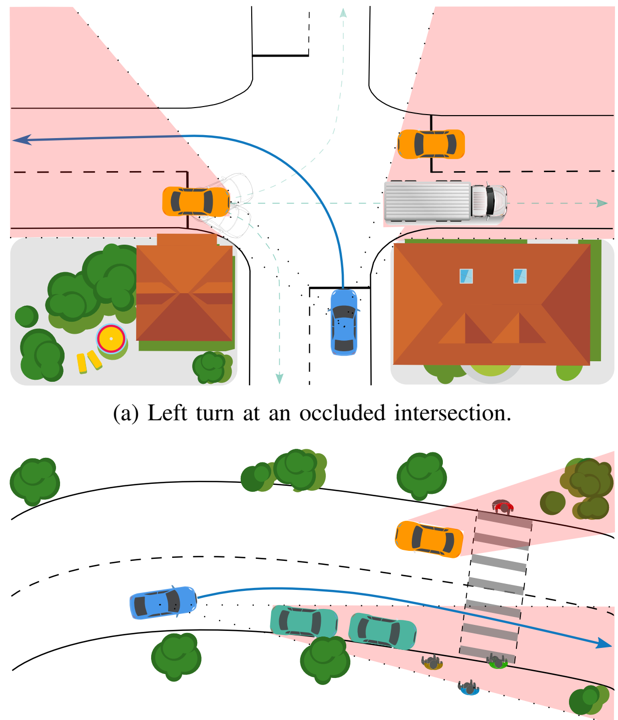

(a) Left turn at an occluded intersection.

(b) Occluded road section and crosswalk.

Fig. 1: Two typical driving scenarios in an urban area with occlusions due to static obstacles (houses and parked cars) and dynamic obstacles (moving vehicle and truck).

When human drivers pass through these occluded urban areas (Fig. 1), they usually slow down and reason how likely it is that a road user will appear according to the amount of occlusion in the scenario. For example, in an occluded crosswalk shown in Fig. 1b, it is very likely to see pedestrians emerging and crossing the road. In this case, human drivers approach carefully until they have sufficient visibility. A similar driving behavior is also expected when leaving an occluded bus stop or facing an intersection with occluded oncoming lanes. To attain human-like driving behavior, a sophisticated planning algorithm should incorporate context information and road topology to estimate the probability of the appearance of road users and their possible future interaction with the ego vehicle.

[1]Chi Zhang, [2]Florian Steinhauser are with ZF Friedrichshafen AG, Friedrichshafen, Germany. (chi.zhang@zf.com, florian.steinhauser@zf.com)

[3]Gereon Hinz, [4]Alois Knoll are with Chair of Robotics, Artificial Intelligence and Real-time Systems, Technical University of Munich, Germany. (gereon.hinz@tum.de, knoll@mytum.de)

In this work, we extend a partially observable Markov decision process (POMDP)-based behavior planner by combining context-aware phantom road users to handle multiple occlusion scenarios in an urban environment. First, we identify risky occlusion areas along the ego navigation path based on map information. These areas can be intersections, bus stops, and crosswalks. Based on identified occlusion areas, we further generate potential hidden road users and infer their movements, such as vehicles in driving lanes or pedestrians at crosswalks whose intention is to cross the road. We then introduce the probability of phantom road users appearing outside of the occluded area. The appearance probability consists of an area-specific part that considers map information and a dynamic part that captures the change in the ego vehicle field of view (FoV) in future time steps. Finally, we combine the context-aware phantom road users with a probabilistic POMDP to plan safe maneuvers considering the risk due to phantom objects.

In summary, the contributions of this study are as follows:

- the extension of the phantom road user concept to include pedestrians to improve the occlusion scenario coverage of a POMDP-based decision-making system,
- a context-based appearance probability method which easily incorporates context information and road topology in a POMDP to plan deadlock-free and comfortable driving behavior in the presence of heavy occlusions,
- the evaluation of the algorithm for handling multiple complex occlusion scenarios, including occluded crosswalks, bus stops (marginally studied but important for autonomous shuttle buses), and intersections in urban environments.

The rest of this paper is structured as follows. The related work on occlusion-aware behavior planning approaches is discussed in Section II. In Section III, the concept of this work is described in detail. The POMDP implementation is presented in Section IV. Evaluation results are shown in Section V. Finally, in Section VI conclusions are drawn and the direction of future work is discussed.

II. RELATED WORK

The purpose of occlusion-aware behavior planning is to make the ego vehicle drive efficiently without overcautiousness while reacting safely if a road user suddenly appears from an occluded area. Recently, several techniques have been proposed to address sensor limitations and occlusions.

Some researchers performed reachability-based analysis with worst-case assumptions to consider the risks due to potential traffic participants in occluded areas. The general idea is to use a set of states to represent all possible configurations that a vehicle could reach. [1] computed the longitudinal speed profile for an AV with path-velocity decomposition. Occluded vehicles were incorporated as dynamic constraints generated with worst-case scenarios by modeling virtual vehicles with infinite length and maximum speed. Similarly, [2] performed reachability analysis to analyze the safety of passing through a potential conflict area with occluded vehicles. The occurrence probability was used as a threshold. If it was sufficiently low, the uncomfortable emergency stop was also acceptable. [3] applied a set of particles to represent potential configurations of occluded vehicles in a region of interest and analyzed the collision risk with a predicted visibility range in future time steps. Then the predicted visibility risk was combined with a cost-based planner to plan the acceleration for the AV. Reachability analysis can prove safety but also results in conservative driving behavior in some special cases with very limited visibility [1]. In this case, carefully advancing into a conflict zone is necessary to gather more information instead of freezing.

Contrary to worst-case assumptions, learning-based methods focus on automatically learning complex driving strategies from data without hand-coded rules. Reference [4] applied a deep Q-network (DQN) to learn the policy in an unsignaled intersection with static occlusion. The state space was described as an occupancy grid map to separate drivable regions and obstacles. [5] focused on safe reinforcement learning using a model-checker to identify safe actions from the action space. Potential incoming road users were modeled via additional state variables. [6] handled occluded intersections using a risk-aware DQN, which incorporated risk evaluation within the reward function instead of only considering collisions. A vehicle with maximum allowed velocity was assumed if the intersection was occluded. These approaches give ideas for applying learning-based methods to generate driving behaviors. However, an approach providing both safe and robust driving behavior in the presence of unseen scenarios or corner cases when observed data varies slightly from training data still needs to be further investigated.

Another planning category focused on probabilistic models, which integrate uncertainty due to visibility limitation in the POMDP model. Previous works in this area can be classified into two groups depending on whether the policies are solved offline or online. The offline approaches calculate the approximated optimal policies over the entire state space for an arbitrary initial belief. [7] combined a rule-based autonomous emergency braking system with POMDP to obtain less conservative driving behavior for a pedestrian collision system under sensor occlusion. In [8], static occlusion in a crosswalk and at a simple T-junction was considered in an offline POMDP. A scene decomposition method was introduced by treating each road user independently to improve the scalability of the POMDP with multiple road users. However, the authors did not consider occlusions due to dynamic moving objects. In addition, a state space model encompassing road users, map geometries, and traffic rules must be solved before using POMDP online, which makes the POMDP model more difficult to solve and restricts its application to a highly dynamic urban environment. By contrast, online POMDP solves the problem online by constructing belief trees with a reachable set of states

TABLE I: Comparison of existing occlusion-aware planning methods with the proposed approach.

Source	Static obstacles	Dynamic obstacles	Vehicles	Pedestrians	Intersections	Crosswalks	Bus stops
[1]	✓		✓		✓		
[2]	✓	✓	✓		✓		
[3]	✓	✓	✓		✓		
[4]	✓		✓		✓		
[5]	✓		✓	✓	✓	✓	
[6]	✓		✓		✓		
[7]	✓			✓		✓	
[8]	✓		✓	✓	✓	✓	
[9]	✓	✓	✓		✓		
[10]	✓	✓	✓		✓		
[11]	✓	✓	✓		✓		
Ours	✓	✓	✓	✓	✓	✓	✓

for the current belief within a limited time horizon. In [9] and [10], phantom vehicles were placed on the edge of invisible areas at unsignalized intersections and treated as real vehicles in the planning phase. [11] considered traffic density and modeled the occurrence probability for phantom vehicles. The abovementioned approaches can handle occluded scenarios arising from static and dynamic objects but are limited to unsignalized intersections.

The state-of-the-art solutions showed the progress of the decision-making system under visibility limitations. However, sophisticated occlusion scenarios, such as bus stops and crosswalks where lots of vulnerable road users could suddenly emerge, are less discussed but important for autonomous shuttle buses (see comparison listed in TABLE I). In this work, we extend the POMDP-based behavior planner by introducing context-based appearance probability for phantom vehicles and pedestrians to handle multiple complex occluded scenarios.

III. OCCLUSION-AWARE ONLINE PLANNING

This study focuses on the high-level longitudinal behavior decision-making system for the AV. By using the path-velocity decomposition concept [12], the action generated from the POMDP planner is applied to control an ego vehicle along a planned driving path in the longitudinal direction. This section first introduces the general theory of POMDP. Then, the extended observation model of POMDP to generate context-aware phantom road users is described. Finally, the combination of probabilistic appearance of phantom road users within the POMDP transition model is explained.

A. POMDP Preliminaries

A POMDP is a probabilistic method that models the sequential decision process of a system (often denoted as agent) under uncertain conditions. A POMDP is defined by the tuple $(S, A, O, T, Z, R, \gamma)$ where S, A, O represent the state, action, and observation spaces, respectively. The transition model T is a conditional probability function $T(s,a,s') = P(s' \mid s,a)$ modeling the probability of a system transition from the state $s \in S$ to the state $s' \in S$ when action $a \in A$ is executed. Similarly, the observation model Z is a conditional probability function $Z(o,a,s') = P(o \mid s',\ a)$ describing the probability of receiving observation $o \in O$ after taking action $a \in A$ and transitioning to state $s' \in S$. The reward $R(s,\ a)$ is the immediate reward generated by performing the action $a \in A$ from the state $s \in S$. Finally, a factor $\gamma \in [0,1)$ discounts future rewards [13].

In a partially observable environment, the agent has only partial knowledge of the system state. Hence, a belief state $b(s)$ is maintained to reflect its internal knowledge of the system and estimates the true state. The policy $\pi : B \to A$ is a mapping from a belief $b \in B$ to an action $a \in A$. Therefore, the solution to a POMDP problem is an optimal policy π^* that maximizes the expectation of accumulated reward over time:

$$\pi^* = \arg\max_{\pi} E\left[\sum_{t=0}^{\infty} \gamma^t R(s_t, \pi(b_t)) \mid b_0, \pi\right]. \tag{1}$$

B. State Space

The state space contains an ego vehicle, surrounding traffic participants, and phantom objects, which includes vehicles and pedestrians. However, road context information is not modeled in the state space since it does not change and can be accessed during the planning cycle. The ego state includes the Cartesian position $(x,\ y)$, orientation θ, velocity v, and planned ego route r, which is obtained either from map geometry or from the path planner after setting a driving destination. It is defined as $s_{ego} = [x_{ego}, y_{ego}, \theta_{ego}, v_{ego}, r_{ego}]^T$. The state of other considered objects s_i is also described in the same manner. The intention of other objects r_i is the partially observable variable that can only be inferred from observation. The state $s \in S$ can be denoted as a joint state vector: $S = [s_{ego}, s_1, s_2, ..., s_N]^T$, where s_{ego} is the state of ego vehicle. The states s_i with $i \in \{1, 2, \ldots, N\}$ represent the dynamic objects. These can be either surrounding traffic participants or phantom objects placed in the occluded region of interest.

C. Action Space

The design goal of the action space is to use as few actions as possible to represent a wide variety of behaviors, such as slowing down, stopping in front of a crosswalk, and accelerating to drive through a junction. To achieve that, we apply a discrete set of acceleration values to represent the maneuvers $A = \{+1.5\ m/s^2, 0\ m/s^2, -1.5\ m/s^2\}$: acceleration, keep velocity, deceleration, respectively.

D. Observation Space

Similarly, the observation space $o \in O$ is a combination of each observation including the ego vehicle, surrounding traffic participants, and phantom objects, which can be denoted as follows: $O = [o_{ego}, o_1, o_2, \ldots, o_N]^T$, with $o_{ego} =$

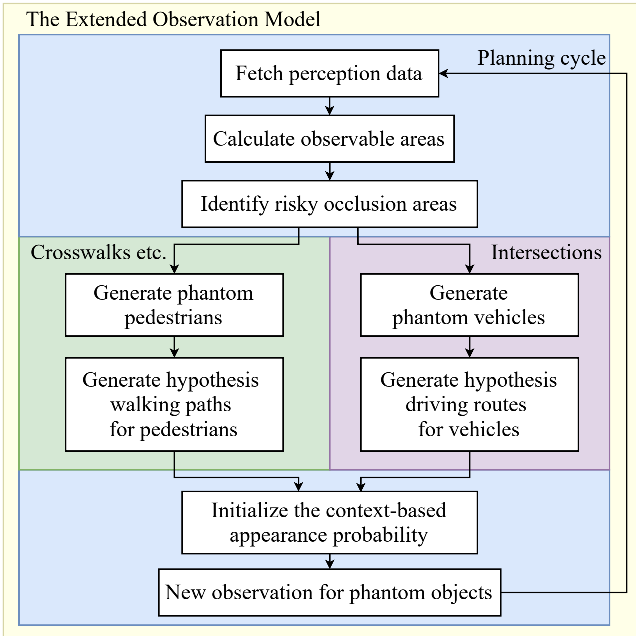

Fig. 2: The extended observation model for belief update of phantom objects.

$[x_{ego}, y_{ego}, \theta_{ego}, v_{ego}, r_{ego}]^T$ being the observation of ego vehicle and $o_i = [x_i, y_i, \theta_i, v_i, r_i]^T$ with $i \in \{1, 2, \ldots, N\}$ representing the observation of all traffic participants and phantom objects.

E. Observation Model

The observation model is the representation of the uncertain sensor measurements. It describes the partial view of the ego vehicle in the state space. The observation model $Z(o, a, s') = P(o \mid s', a)$ is modeled as the conditional probability:

$$Z\left(o, a, s'\right) = \mathrm{P}\left(o \mid s', a\right) = \mathrm{P}\left(o_{ego} \mid s'_{ego}, a_{ego}\right) \prod_{i=1}^{N} \mathrm{P}\left(o_i \mid s'_i, a_i\right). \tag{2}$$

Thus, the observation of the ego vehicle and other agents can be treated individually.

F. Observation Model of Real Objects

The observation of the ego o_{ego} can be generated directly from measurements. For other real traffic participants, their position, orientation, and speed can also be observed directly. Their intentions are inferred and updated by the ego vehicle's prediction module whenever new measurements are received.

G. Extended Observation Model of Phantom Objects

Phantom objects are defined as objects that cannot be observed. We extend the observation model for the generation of phantom road users (see Fig. 2). Since phantom objects can suddenly emerge and lead to potentially dangerous situations, their possible occurrence from the unobservable area needs to be inferred. Based on map information, we generate phantom vehicles with corresponding coordinates, speeds, and potential intentions regarding which lane to take for the occluded intersections (see Fig. 3a). As shown in Fig. 3d, phantom pedestrians, whose goal is to cross the road, are generated in occluded crosswalk areas, bus stops etc.

1) Generation of phantom road users: The first step is to calculate observable areas based on static and dynamic obstacle information obtained from the perception module. The observable area is first initialized with the maximum range of the perception system. Then, the static and dynamic obstacles limit this observable area represented by a polygon. Next, we search for potential risk areas in the map along the upcoming navigation path of the ego vehicle from the occlusion catalog, which contains risk occlusion scenarios that need to be handled. For example, in a structured intersection without a traffic light, we choose occluded lanes with higher priority than ours. Moreover, we consider unobservable areas on or near crosswalks and bus stops, where pedestrians have a higher appearance probability. After selecting the occluded risk areas, the phantom traffic participants are placed on each edge of the FoV. The lengths of the phantom objects are defined to be infinite, which enables us to represent a set of reachable states using only one configuration of phantom objects. Their potential paths also need to be determined. All following lanes for phantom vehicles in the intersection are extracted from the map as candidate paths. For the phantom pedestrians, pseudo priority walking paths are generated. The walking paths consist of waypoints starting at the occlusion edge and point to the other side of the road.

2) Context-based appearance probability: The next question that needs to be addressed is the probability of a phantom object entering the observable area. Assuming that the phantom object always appears at each planning cycle would cause the ego vehicle to drive overcautiously. In some cases, this assumption could even block the ego vehicle and lead to a "freezing state" without any further movement. Our idea is to incorporate context information into the appearance probability:

$$P_a(d, u) = \min\left(\left(P_{env}(d) + P_{FoV}(u)\right),\ 1\right). \tag{3}$$

The min operator is employed to guarantee an appearance probability $P_a(d, u) \leqslant 1$. The context-based appearance probability $P_a(d, u)$ consists of two parts. The first part $P_{env}(d)$ represents the environmental context (see Eq. 4). We define K_{env} as an initial environmental probability to reflect the phenomenon that the appearance of road users depends on where the occlusion occurs. For example, the probability of pedestrians appearing at crosswalks is greater than that on ordinary roads. Similarly, pedestrians are more likely to appear around bus stops and school areas. If the occluded area is far away from a crosswalk, the appearance probability is small. Hence, a distance threshold D_s is introduced. The distance d from the start point of the phantom object to this risky region is included in a discount factor, which means

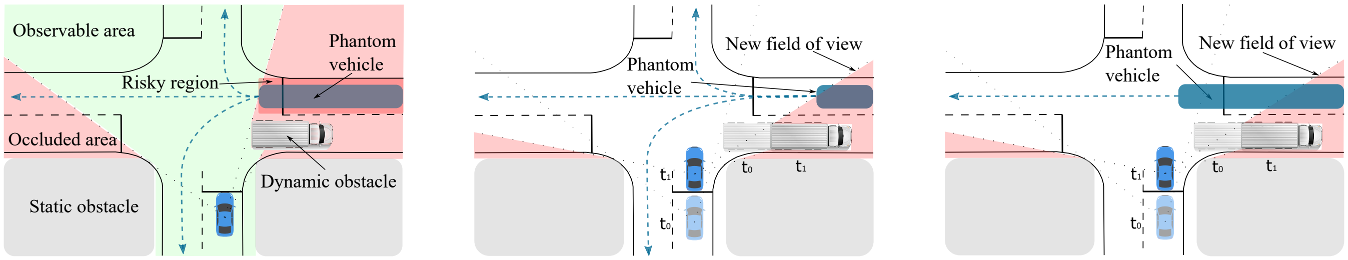

(a) Observation model of phantom vehicle in intersection.

(b) Phantom vehicle stays in occlusion area.

(c) Phantom vehicle drives out of occlusion area.

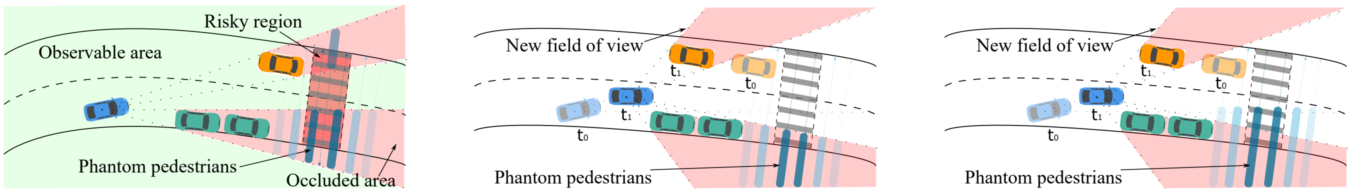

(d) Observation model of phantom pedestrians in crosswalk.

(e) Phantom pedestrians stay in occlusion area.

(f) Phantom pedestrians walk out of occlusion area.

Fig. 3: The observation model and the transition model of phantom vehicles and pedestrians. The blue car is the ego vehicle. The yellow car and white truck are other moving traffic participants. Green cars represent parked cars. For phantom pedestrians, dark blue means high appearance probability.

that the area-specific appearance probability is no longer considered beyond a certain distance threshold D_s.

$$P_{env}(d) = \max\left(\left(K_{env}\frac{D_s - d}{D_s}\right),\ 0\right). \tag{4}$$

The second part of the probability $P_{FoV}(u)$ describes the probability of phantom objects' appearance due to a change in the FoV. The FoV can change during forward simulation in each time step. If the FoV is enlarged, the chance that the ego vehicle observes a phantom object previously hidden by obstacles increases. Inspired by [11], we assume that phantom vehicles and pedestrians are normally distributed in the occluded lane or hypothesis path. We define L as the length within which we expect to observe exactly one phantom object. With the increase in the FoV of u meters, the probability for at least one phantom object to appear is $\frac{u}{L}$. In case the FoV decreases or remains constant, no additional probability of observing phantoms is assigned.

$$P_{FoV}(u) = \begin{cases} 0 & ,\ \text{for } u <= 0 \\ \frac{u}{L} & ,\ \text{for } u > 0 \wedge u < L \\ 1 & ,\ \text{for } u >= L. \end{cases} \tag{5}$$

Based on the defined appearance probability $P_a(d,u)$, we can sample whether a phantom road user comes out of the occlusion area.

H. Extended Transition Model

The system's stochastic dynamics are determined by the transition model $T(s,a,s')$. According to Bayes's rule, the transition model can be described by the probabilistic transition in Eq. 6, which means the transition model for each agent can be calculated individually.

$$\mathrm{T}(s,a,s') = \mathrm{P}(s' \mid s,a) = \mathrm{P}(s'_{ego} \mid s_{ego}, a_{ego}) \prod_{i=1}^{N} \mathrm{P}(s'_i \mid s_i, a_i). \tag{6}$$

1) Ego vehicle: A point mass dynamic model is applied. Since the ego vehicle's mission path is known, the transition of the ego vehicle is determined in Frenet coordinates [14] by applying a distance progress l along the mission path:

$$\begin{bmatrix} l' \\ v' \\ r' \end{bmatrix} = \begin{bmatrix} 1 & \Delta t & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} l \\ v \\ r \end{bmatrix} + \begin{bmatrix} \frac{1}{2}\Delta t^2 \\ \Delta t \\ 0 \end{bmatrix} a. \tag{7}$$

The new ego state s'_{ego} in Cartesian coordinates is obtained by the transforming $[l', v', r']^T$ from Frenet coordinates.

2) Other objects: The state transition model for other objects is also referred to as the prediction model in the literature. The prediction model estimates other traffic participants' intentions and future states based on dynamics and context information. Through the interface with the prediction module, we can obtain the prediction and uncertainty of the next state given the current state and the intention of other objects. Thus, sophisticated prediction models can be used for different road users. Since this is not the focus of this study, we apply a simple constant velocity model to update the state of other objects along their path.

3) Phantom objects: When updating the state transition of phantom objects over time step Δt, we first sample according to the context-based appearance probability $P_a(d,u)$. If the sample result is zero, the phantom object is placed at the updated edge of the FoV, as shown in Figs. 3b and 3e. Otherwise, when the sample result is one, its position l'_l is moved forward using a constant velocity model: $l_l' = l_l + v_{\max} \cdot \Delta t$ (see Figs. 3c and 3f), with $v_{\max}$ being the maximum velocity allowed in the driving lane. For the pedestrian, we choose $v_{\max} = 1.25\ m/s$ to represent a normal walking speed. Once the phantom object is outside of the occluded area, it will only move forward along the path within a planning cycle.

I. Reward Model

The reward function is a crucial factor in designing the ego vehicle's behavior which satisfies several objectives, such as safety, efficiency, and comfort. We encode these objectives mathematically within the reward function $R(s,\ a)$:

$$R = R_{collision\_real} + R_{collision\_phantom} + R_{speed} + R_{comfort}. \quad (8)$$

A number of simulations were performed in order to find suitable weights for the reward function. Safety can be modeled by assigning a large negative reward $R_{collision\_real} = -100000$ if the ego vehicle collides with other road users. We assign a different penalty in the case of a collision with a phantom object $R_{collision\_phantom} = -10000$.

Speed reward is used to encourage the ego vehicle to drive according to the desired velocity $v_{desired}$ on the driving lane as far as possible while not exceeding it:

$$R_{speed} = \begin{cases} -200 \cdot (v_{desired} - v_0), & \text{if } v_{desired} \geq v_0 \\ -2000 \cdot |v_{desired} - v_0|, & \text{otherwise.} \end{cases} \quad (9)$$

Comfort is considered by penalizing the changing accelerations: $R_{comfort} = -300 \cdot a^2$.

IV. IMPLEMENTATION

Urban traffic is a partially observable and highly dynamic environment. Therefore, the applied POMDP solving algorithm should be able to handle large continuous state spaces and provide at least near-optimal solutions in real-time. This study applies the TAPIR toolkit [15] to solve the previously described POMDP model.

A. Online Belief Tree Construction

TAPIR builds a belief tree and searches for sequential actions that maximize the accumulation of reward. The fundamental strategy for constructing a belief tree and improving the policy is to sample the episodes. To sample a new episode, TAPIR begins by sampling a set of unweighted particle states from an initial belief b_0, which approximate the environment's initial state. A belief tree τ is built from this belief state by sampling further episodes representing a set of historical sampled belief, action, reward, and observation tuples. As a result, TAPIR only considers reachable states from the current belief, which constrains the search space dramatically and is capable of delivering a near-optimal policy for a given search time.

TABLE II: Parameters applied in evaluation.

Parameter	Value	Parameter	Value
Planning Frequency	2 Hz	Planning Horizon	10 s
Discount Factor γ	0.95	Maximal Tree Depth	10
Scenario A: $V_{pedestrian1}$	1.5 m/s	$V_{pedestrian2}$	1.0 m/s
Scenario B: $V_{pedestrian1}$	0.8 m/s		
Scenario C: V_{truck}	5.0 m/s	$V_{vehicle}$	6.0 m/s

B. Varying Planning Time Steps

To reduce the total number of planning steps and to consider more accurate predictions closer to the current observation, we use a search strategy with varying time steps. The time steps are defined as follows: $[\Delta t_0 ... \Delta t_3, \Delta t_4 ... \Delta t_7, \Delta t_8, \Delta t_9] = [0.5\ s ... 0.5\ s, 1.0\ s ... 1.0\ s, 2.0\ s, 2.0\ s]$.

C. Tunnel Effect Avoidance

At each step, the collision status between the ego vehicle and other objects needs to be detected. However, if the time step is too large or if the relative speed is too high, the two objects may miss each other, resulting in a possible collision not being detected. Therefore, inspired by [16], we use the linearity assumption between each step to detect whether two moving objects cross each other. For phantom objects, the collision is detected by checking whether line segments intersect.

V. EVALUATION

In this section, we evaluate our approach in a proprietary simulator under various challenging occlusion scenarios, including crosswalks, bus stops, and intersections. Additional scenarios are presented in the supplementary video[1]. To eliminate the influence of other road users' intelligent behavior in the evaluation, we control them via predefined behaviors that do not consider collision avoidance. The velocity of all road users in the evaluation is chosen to compare different planning behaviors. The parameters are listed in TABLE II.

We compare our approach's performance against two other strategies. Our generic occlusion-aware POMDP behavior planner is denoted as GO-POMDP. V2X-POMDP represents a ground truth planner, which has access to all available environmental information. Finally, we set up a POMDP planner (WO-POMDP) with the worst-case assumption that phantom objects will always appear from occluded areas with maximal allowed velocity. The worst-case assumption has been widely applied in other studies [1], [9]. Thus, we treated it as a baseline approach for comparison against our approach.

A. Occlusion in Crosswalk

The first scenario (see Fig. 4) is a crosswalk that is partially occluded by a parked vehicle. Two pedestrians in the occluded area will cross the road when the ego vehicle

[1] Video: https://github.com/GitChiZhang/GO-POMDP

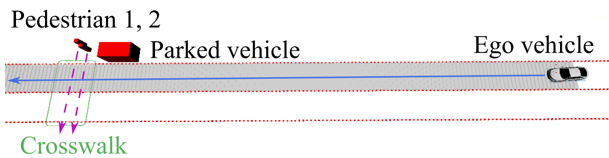

Fig. 4: Ego vehicle driving through an occluded crosswalk. Two pedestrians are about to cross the crosswalk, but their view is blocked by the parked car.

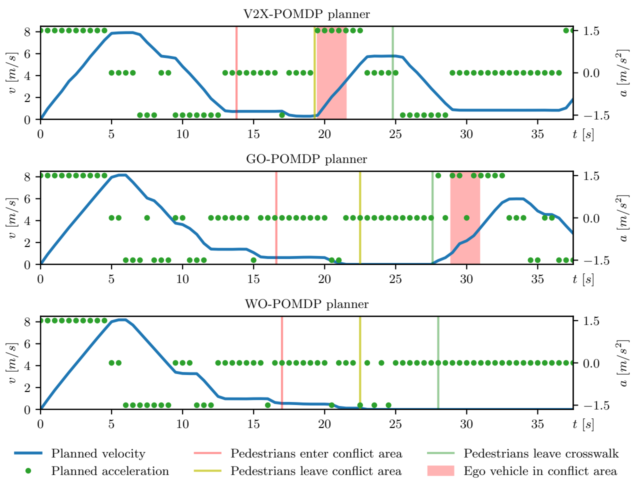

Fig. 5: Planned velocity and acceleration profiles for crosswalk scenario with two occluded pedestrians.

approaches near the crosswalk. The result in Fig. 5 shows that V2X-POMDP adjusts its velocity when approaching the crosswalk because it fully observes the pedestrians. We can also see that V2X-POMDP behaves aggressively as it starts to accelerate as soon as the pedestrians leave the conflict area, i.e., walk toward the other lane of the road. Similarly, GO-POMDP reduces its velocity to approach the crosswalk carefully due to occlusion awareness. Because of limited sight in the crossing area, the ego vehicle moves forward at a very low speed. After the pedestrians appear, it first maintains its speed since it still has sufficient safe distance to the pedestrians. At time $t = 20$ s, the ego vehicle comes to a halt and lets the pedestrians pass first. The yellow and green lines of GO-POMDP in Fig. 5 indicate that the ego vehicle is waiting for pedestrians to leave the road completely. This is because the moving pedestrians may block the view of the ego vehicle to observe pedestrians coming from the other direction, which has been considered as a risk in GO-POMDP. WO-POMDP has similar driving behavior when approaching the crosswalk. However, due to the limited visibility, the ego vehicle will not continue to drive under the worst-case assumption that 100% of pedestrians will appear from the occluded area. This "freezing state" has also been reported in the previous study [1].

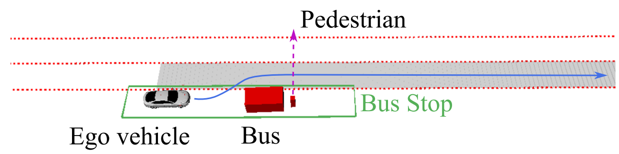

Fig. 6: Ego vehicle is leaving the bus stop while a pedestrian is about to cross the road in front of a bus.

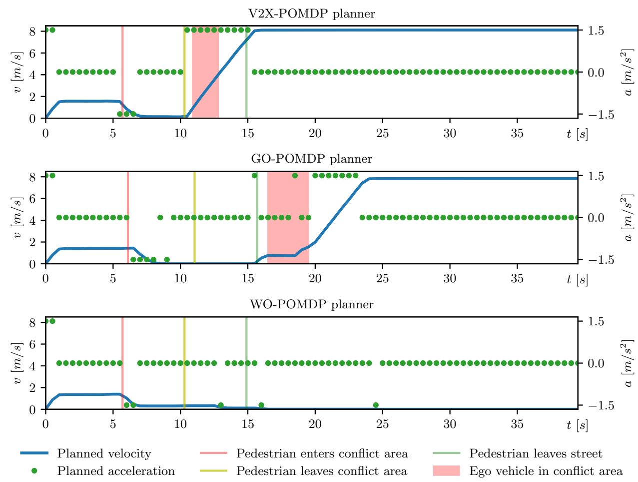

Fig. 7: Planned velocity and acceleration profiles for bus stop scenario with an occluded pedestrian.

B. Occlusion in Bus Stop

The second occluded scenario (Fig. 6) is at a bus stop. While the ego vehicle is about to leave the bus stop, a passenger wants to cross the road in front of the bus after exiting it. This is also a challenging scenario as the ego vehicle can't detect this pedestrian before it appears on the ego path. The results in Fig. 7 show that V2X-POMDP drives faster than the other two planners through the bus stop since it has full knowledge of the occluded road user. GO-POMDP performs very well in terms of avoiding a collision with the pedestrian. After the pedestrian leaves the high-risk area, GO-POMDP continues to drive with a low velocity to increase the visibility while driving through this area. In contrast, WO-POMDP once again leads to a deadlock situation due to overcautiousness.

C. Occlusion in Unsignalized Intersection

Finally, we demonstrate our approach at an unsignalized intersection with a left-turn maneuver under dynamic occlusion (see Fig. 8). A moving truck arrives at the intersection before the ego vehicle and prevents the ego vehicle from observing another vehicle that has priority. Fig. 9 shows that V2X-POMDP waits for the vehicle on the prioritized lane after the moving truck has crossed the intersection. It can also be seen that GO-POMDP performed nearly as well as V2X-POMDP. From $t = 10$ to 13 s, it slows down and performs creeping behavior to observe whether there is a vehicle on the priority lane occluded by the truck.

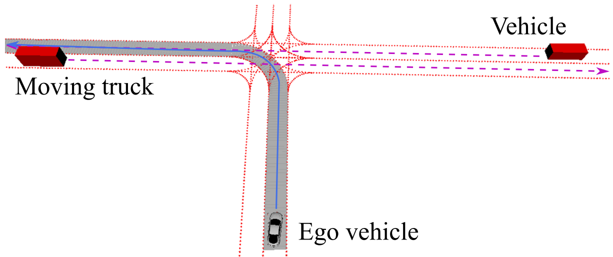

Fig. 8: Left turn in an unsignalized intersection scenario with dynamic occlusion due to a moving truck.

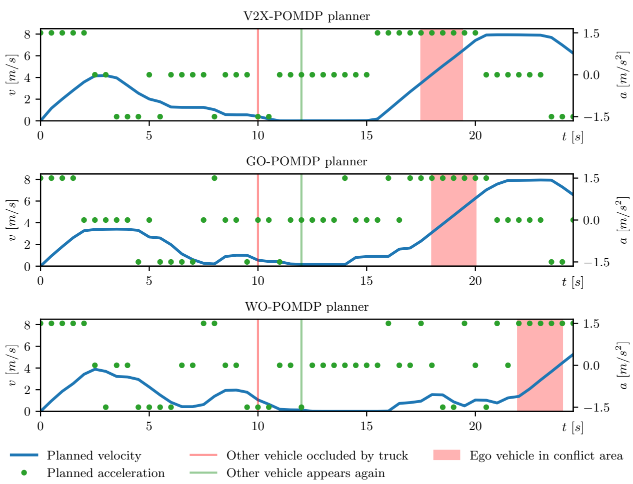

Fig. 9: Planned velocity and acceleration profiles for unsignalized intersection scenario.

Once the other vehicle appears and has left the conflict area, GO-POMDP accelerates slightly and enters the intersection after having sufficient visible area. WO-POMDP behaves more cautiously than our approach.

VI. CONCLUSION AND FUTURE WORK

In this study, we improve the occlusion scenario coverage with a POMDP-based behavior planner for driving in urban environments under occlusions. Our approach handles various occlusion situations due to static or dynamic obstacles for different occluded road users. Our main contribution is to combine the uncertainty consideration of POMDP and context-aware phantom object modeling with appearance probability, which enables the effective use of map information. Evaluation results show that the ego vehicle can safely drive through challenging occlusion scenarios, such as crosswalks, bus stops, and intersections. The ego vehicle performs comfortable driving behavior under occlusions, i.e., creep forward into the conflict area to increase visibility with the knowledge that other road users could suddenly emerge. Owing to the effective modeling of the context-based appearance probability, our approach does not cause deadlocks as the worst-case assumption approach would in heavy occlusion situations. In the future, we plan to evaluate our approach using more simulations and real-world experiments. Furthermore, in some scenarios, the visibility of other road users is also limited. In such cases they may not yield even though the ego vehicle has a higher priority. We will extend our approach to predict the visibility changes of other potential road users and consider their potential interaction with the ego vehicle. In addition, we plan to design automatic corner cases generation to find optimal weights for the reward function.

REFERENCES

[1] R. Poncelet, A. Verroust-Blondet, and F. Nashashibi, "Safe geometric speed planning approach for autonomous driving through occluded intersections," in *2020 16th International Conference on Control, Automation, Robotics and Vision (ICARCV)*. IEEE, 2020, pp. 393–399.

[2] M. Naumann, H. Konigshof, M. Lauer, and C. Stiller, "Safe but not overcautious motion planning under occlusions and limited sensor range," in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 140–145.

[3] L. Wang, C. F. Lopez, and C. Stiller, "Generating efficient behaviour with predictive visibility risk for scenarios with occlusions," in *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2020, pp. 1–7.

[4] D. Isele, R. Rahimi, A. Cosgun, K. Subramanian, and K. Fujimura, "Navigating occluded intersections with autonomous vehicles using deep reinforcement learning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 2034–2039.

[5] M. Bouton, A. Nakhaei, K. Fujimura, and M. J. Kochenderfer, "Safe reinforcement learning with scene decomposition for navigating complex urban environments," in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 1469–1476.

[6] D. Kamran, C. F. Lopez, M. Lauer, and C. Stiller, "Risk-aware high-level decisions for automated driving at occluded intersections with reinforcement learning," in *2020 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2020, pp. 1205–1212.

[7] M. Schratter, M. Bouton, M. J. Kochenderfer, and D. Watzenig, "Pedestrian collision avoidance system for scenarios with occlusions," in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 1054–1060.

[8] M. Bouton, A. Nakhaei, K. Fujimura, and M. J. Kochenderfer, "Scalable decision making with sensor occlusions for autonomous driving," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 2076–2081.

[9] X. Lin, J. Zhang, J. Shang, Y. Wang, H. Yu, and X. Zhang, "Decision making through occluded intersections for autonomous driving," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2019, pp. 2449–2455.

[10] P. Schörner, L. Töttel, J. Doll, and J. M. Zöllner, "Predictive trajectory planning in situations with hidden road users using partially observable markov decision processes," in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 2299–2306.

[11] C. Hubmann, N. Quetschlich, J. Schulz, J. Bernhard, D. Althoff, and C. Stiller, "A pomdp maneuver planner for occlusions in urban scenarios," in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 2172–2179.

[12] K. Kant and S. W. Zucker, "Toward efficient trajectory planning: The path-velocity decomposition," *The international journal of robotics research*, vol. 5, no. 3, pp. 72–89, 1986.

[13] M. J. Kochenderfer, *Decision making under uncertainty: theory and application*. MIT press, 2015.

[14] M. Werling, J. Ziegler, S. Kammel, and S. Thrun, "Optimal trajectory generation for dynamic street scenarios in a frenet frame," in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 987–993.

[15] D. Klimenko, J. Song, and H. Kurniawati, "Tapir: A software toolkit for approximating and adapting pomdp solutions online," in *Proceedings of the Australasian Conference on Robotics and Automation, Melbourne, Australia*, vol. 24, 2014.

[16] S. Brechtel, T. Gindele, and R. Dillmann, "Probabilistic decision-making under uncertainty for autonomous driving using continuous pomdps," in *17th international IEEE conference on intelligent transportation systems (ITSC)*. IEEE, 2014, pp. 392–399.