Looking for Trouble: Informative Planning for Safe Trajectories with Occlusions

Barry Gilhuly, Armin Sadeghi, Peyman Yedmellat, Kasra Rezaee, Stephen L. Smith

***Abstract*— Planning a safe trajectory for an ego vehicle through an environment with occluded regions is a challenging task. Existing methods use some combination of metrics to evaluate a trajectory, either taking a worst case view or allowing for some probabilistic estimate, to eliminate or minimize the risk of collision respectively. Typically, these approaches assume occluded regions of the environment are unsafe and must be avoided, resulting in overly conservative trajectories—particularly when there are no hidden risks present.**

We propose a local trajectory planning algorithm which generates safe trajectories that maximize observations on uncertain regions. In particular, we seek to gain information on occluded areas that are most likely to pose a risk to the ego vehicle on its future path. Calculating the information gain is a computationally complex problem; our method approximates the maximum information gain and results in vehicle motion that remains safe but is less conservative than state-of-the-art approaches. We evaluate the performance of the proposed method within the CARLA simulator in different scenarios.

I. Introduction

The safety of an ego vehicle depends on its understanding of the risk in the world around it. A key challenge is the presence of dynamic and static objects, which may result in occlusions in the sensor information received by the ego vehicle. These occlusions hide regions of the environment and prevent the detection of obstacles in the vehicle's intended trajectory. The safe action when presented with an occlusion is to anticipate some amount of risk due to the potential of hidden actors or other hazards [1]. This anticipation of risk may result in overly conservative responses, which slow the ego vehicle excessively or even bring it to a complete halt in anticipation of a perceived threat.

A trajectory with no occlusions has fewer unknowns and therefore less uncertainty in the environment. Less uncertainty results in a safer trajectory, or allows the vehicle to proceed at a faster rate while maintaining the same safety constraints. One method of reducing occlusions is to generate trajectories that are allowed to deviate from the reference center line of the lane [2]. With slight deviations, more regions of the environment may be revealed.

Road networks are complicated because the paths of multiple actors cross and possibly conflict. Occlusions in this environment are complex, with multiple obstructions, both static and dynamic, introducing uncertainty [3]–[6]. Because of this complexity, it may not be possible to eliminate occlusions by adjusting the path of the trajectory; doing so may merely re-position the occlusions to conceal a different region. The key then is to determine which occlusions are a greater source of planning uncertainty and preferentially minimize those first.

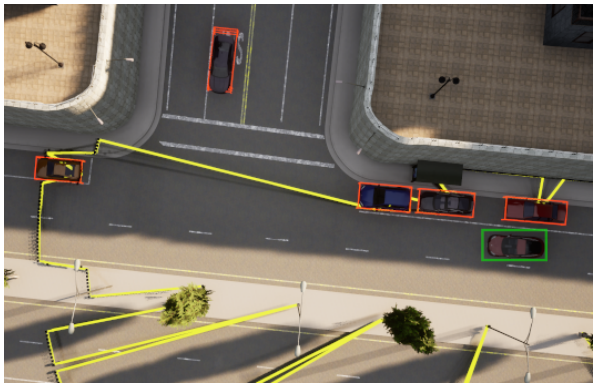

Fig. 1: Ego vehicle (green) approaching an intersection with obstructed vision. Current visibility is outlined in yellow.

In this paper, we address the problem of evaluating trajectories for the ego vehicle that minimize environmental uncertainty by maximizing observations on the most critical occluded regions. Maximizing critical observations, referred to as maximizing a trajectory's information gain, identifies those trajectories that make observations of regions of the environment that pose a higher risk to the vehicle in its future time steps. For example, consider Figure 1 where an ego vehicle is approaching an intersection while its view is obstructed. Instead of slowing down in anticipation of a collision risk with an actor "hiding" in the occluded region, our proposed informative motion planner deviates from the center line to gather information on the state of the occluded region earlier in its trajectory. Instead of fixing such a behaviour, we allow it to emerge naturally as a result of our information planning framework. Efficient calculation of the true information gain is extremely difficult due to the complex interaction of trajectories, other actors and occlusions in the environment. We present here an approximation method that allows for fast computation. Improving the understanding of the occluded regions allows the ego vehicle to plan less conservative motions while preserving or improving overall safety.

A. Contributions

Our main contributions are threefold. First, we formulate the problem of risk aware motion planning for an ego vehicle,

B. Gilhuly, A. Sadeghi, and S. L. Smith are with the Department of Electrical and Computer Engineering, University of Waterloo, Canada {bgilhuly,a6sadegh,stephen.smith}@uwaterloo.ca

P. Yedmellat and K. Rezaee are with Noah's Arc Lab, Huawei Technologies Canada. {peyman.yadmellat,kasra.rezaee}@huawei.com

This work was supported by Huawei Technologies Canada

introducing information gain as a metric (Section II). Second, we establish a connection between the information gain on the risk of collision and the information gain on a subset of the environment, and propose a method to approximate the information gain for a given trajectory (Section II). Finally, we evaluate the performance of the proposed algorithm within the CARLA Simulator (Section IV).

B. Related Work

The problem of trajectory planning in environments with occlusions is extensively studied. Worst case planners that provide safety guarantees remove potential collision regions from the ego vehicle's workspace [7]–[11]. Many of those potential collision regions are created by occlusions: obscured areas in the environment from which actors could possible emerge [7]. However, worst case planning tends toward an over-estimation of risk. In [4], polygons are computed that approximate the possible regions where actors may be obscured by the environmental occlusions. Authors in [12] present a rigorous set-based analysis of safety, adding phantom vehicles and pedestrians wherever occlusions occur. Different sets of constraints on the dynamics of the actors in the environment are proposed and an over-approximation of the reachability set for each potential actor is proposed to guarantee a safe trajectory. In [13], the authors use occlusion boundary tracking to tighten the bounds on predictions of the speed and location of hidden actors, reducing but not eliminating the overestimation of risk. Another alternative is to use a combination of forward reachability of the ego vehicle combined with sampling the backward reachability of any concealed actor's possible motions to estimate the areas where hidden actors may appear, then excluding those regions from the ego vehicle's workspace [14].

Other methods use a discrete probabilistic representation of the environment (e.g. dynamic grid map) [3], [15]–[18] where each cell provides the probability that the corresponding location is occupied by an actor. The environment around the ego vehicle may be modeled as a dynamic occupancy grid [3], marking cells as being occupied with a certain probability based on observations. The occupancy grid then forms the basis for evaluating potential center line trajectories for safety. Alternatively, the road network may be discretized as a probabilistic grid [17], [19], preserving the structure and allowing application of probabilistic motion predictions conditioned on the expected behaviours of the actors. Conservative motion planning is still required as obscured sections of connected roads, which could be hiding other actors, may lead to situations where the ego vehicle must act slow in anticipation of a possible collision.

In all the aforementioned studies, the ego vehicle takes a passive approach towards minimizing the risk of collision with potential actors in the occluded region. Trajectories are selected by either maintaining the nominal center line trajectory and managing velocity [7], [17] or by selecting a trajectory within the bounds of the reduced workspace [7]. In this paper, we actively minimize the future uncertainty of the risk of collision by selecting safe trajectories which minimize dangerous occluded regions. Our risk evaluation component is based on simple motion prediction; however any method could be used to evaluate trajectory risk, such as in [20] where, given a probabilistic risk measure, the trajectories are tested for potential safety violations.

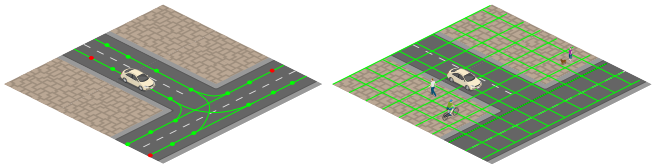

(a) (b)

Fig. 2: The discretization of the environment: a) Road-network G_v capturing the vehicles in the environment, b) Grid G_p capturing the presence of other actors in the environment.

Two closely related problems to maximizing information gain are minimizing occlusions and maximizing the ego vehicle's visibility to other agents. In [21], the robot uses a receding horizon planner to find a trajectory that minimizes the region hidden when approaching a blind corner in a hallway. Similarly, the authors in [22] propose a geometric method to maximize the visibility in overtaking scenarios. In contrast to these two, we propose a method that captures scenarios with multiple sources of sensor obstructions including dynamic actors. We maximize the information gain over multiple occlusions and prioritize observations on regions with high uncertainty. The parallel problem to maximizing visibility is to plan trajectories that minimize travel time in the occlusions of other vehicles [23]. Similar to this work, occlusion-aware planning is used, but for the purpose of proactively ensuring the ego vehicle is both predictable in its behaviour and obscured for as little time as possible.

II. Problem Formulation

Consider an ego vehicle driving in a road network. Alongside the ego vehicle, there are actors operating in the environment such as pedestrians, vehicles, or cyclists. Let q^t be the state of the ego vehicle in the continuous space at time t and let $\Delta t > 0$ be a small time step. A trajectory $\mathcal{T}$ is a sequence of the states of the ego vehicle at times $0, \Delta t, 2\Delta t, \ldots, k\Delta t$, denoted $\mathcal{T} = \langle q^0, q^1, \ldots, q^k \rangle$. We let $\mathcal{T}^{s:e}$ denote the subset of the trajectory on the time interval $[s\Delta t, e\Delta t]$, i.e., $\mathcal{T}^{s:e} = \langle q^s, q^{s+1}, \ldots, q^e \rangle$. A trajectory is *feasible* if it satisfies the motion constraints of the ego vehicle. For simplicity, in the rest of paper we refer to a time $k\Delta t$ simply as time k, and thus with this terminology, all times are integers

The ego vehicle is equipped with sensors and makes observations of the environment. Similar to [19], we let O^{-t} be the set of observations up to time t. The observation at time t, denoted by O^t is a random variable that assigns *occupied*, *free* or *occluded* to each location in the workspace. We define the binary random variable $R(\mathcal{T}^{s:e})$, where $R(\mathcal{T}^{s:e}) = 1$ if the ego vehicle collides with an obstacle on the time interval $[s, e]$ while traversing the trajectory $\mathcal{T}$, and $R(\mathcal{T}^{s:e}) = 0$ otherwise. Given observations O^{-s}, the expected value of

$R(\mathcal{T}^{s:e}|O^{-s})$ represents the probability of collision and it is denoted by $\mathbb{E}(R(\mathcal{T}^{s:e})|O^{-s})$.

Let T be the replan interval, representing the number of time steps before the trajectory is replanned. Then we define the set of collision free states as follows:

Definition 1 (Safe Region)**.** The safe region $SR(t)$ at time t is the set of all configurations q^t of the ego vehicle such that there exists a trajectory $\mathcal{T}_{\text{safe}}$ starting at q^t such that the probability of collision is zero.

Typically, $\mathcal{T}_{\text{safe}}$ is an emergency braking trajectory that the ego vehicle has available. Let Ω be the set of all trajectories $\mathcal{T}$ where $q^t \in SR(t+1)$ for any $t \in [0, T]$. The set Ω represents the set of all trajectories that after one step of executing a trajectory in Ω, there is always a safe trajectory avoiding collisions in the worst-case scenario.

Different measures are proposed in the literature for the quality of trajectory, such as the time it takes to execute the trajectory or the deviation from the nominal trajectory [24]. Let $\mathcal{T}_{\text{nom}}$ be the nominal trajectory of the ego vehicle, which consists of the states on the center lines of the roads, and let NT be the planning horizon for some natural number N. Then we let $J(\mathcal{T}^{0:NT}, \mathcal{T}^{0:NT}_{\text{nom}})$ be the user-defined cost function evaluating the quality of the trajectory over the planning horizon.

Finally, the problem of safe trajectory planning for the ego vehicle is given as follows:

$$\begin{aligned} \min_{\mathcal{T}} \quad & \Bigg(J(\mathcal{T}^{0:NT}, \mathcal{T}^{0:NT}_{\text{nom}}) \\ & + \sum_{n=0}^{N} \rho_n \mathbb{E}\big(\mathbb{E}(R(\mathcal{T}^{nT:(n+1)T})|O^{-nT})\big) \Bigg) \\ \text{such that} \quad & \mathcal{T} \in \Omega, \text{ and } \mathcal{T} \text{ is feasible} \end{aligned} \tag{1}$$

where ρ_n defines relative importance of risk at different planning steps. This optimization is performed at time step $t = 0$, and the observations at future time steps are random variables, therefore by the law of total expectation $\mathbb{E}\big(\mathbb{E}(R(\mathcal{T}^{nT:(n+1)T})|O^{-nT})\big)$ represents the risk of trajectory in time horizon $[nT, (n+1)T]$. The solution to the problem is a feasible trajectory that minimizes the cost and the risk of collision. In Lemma 1, we show that the solution to the problem is a safe trajectory with no collision in the worst case scenario. The planning is performed in a receding horizon manner where after each replan interval T, a new trajectory is computed considering a future horizon of TN steps.

Representation of Environment: The environment is represented by a grid G_p and a directed graph G_v. The pedestrian grid G_p is a simple discretization of the environment and represents the regions occupied by the pedestrians (see Figure 2b). The graph G_v consists of a set of vertices along the center lines of the lanes (see Figure 2a). The center lines are the edges of the graph. Observe that G_v and G_p are different discretizations of the environment. For convenience, we will refer to the vertices on the graph G_v as cells and we let C be the set of all cells in G_v and G_p.

Let X_i^t be a binary random variable representing the state of cell c_i in G_v or G_p, i.e., occupied or free. The probability that a cell c_i is occupied is denoted by $\mathbb{P}(X_i^t)$. We make the simplifying assumption that the random variables are independent. Let $\mathcal{X}$ be the set of all random variables for the cells.

Transition Probabilities: Suppose that the environment contains m actors, where m is unknown to the ego vehicle. Let $\Pr^v_{i,j}(t)$ be the probability that a vehicle actor, located at $c_i \in G_v$, transitions to $c_j \in G_v$ in t time steps. We let $\Pr^p_{i,j}(t)$ represent the t-step transition probabilities of other types of actors residing on G_p. The probability $\Pr^v_{i,j}(t)$ and $\Pr^p_{i,j}(t)$ are proportional to the nominal velocity of the corresponding actor types, the distance between c_i and c_j, and the number of time steps t. Therefore, the set of all cells c_j such that $\Pr^v_{i,j}(t) > 0$ is the *reachable set* of $c_i \in G_v$ with t time steps, denoted by $RS(c_i, t)$. Since the movement of the vehicles is governed by the road network, we estimate the future state of the vehicles in the environment with the locations they can occupy on the road network. Note that the movement of the pedestrians has no implicit structure and therefore the reachable set of a cell $c_i \in G_p$ at time s is the set of all cells within a distance from c_i that a pedestrian can traverse in t time steps, i.e., $\Pr^p_{i,j}(t) > 0$.

Given observations on a cell, we update the probability of the cell using Bayes rule. Let $I(c_i, q^t)$ be an indicator function where $I(c_i, q^t) = 1$ if the rigid body of the ego vehicle at q^t intersects with cell c_i, and $I(c_i, q^t) = 0$ otherwise. Then the risk of trajectory $\mathcal{T}^{s:e}$ is defined as $R(\mathcal{T}^{s:e}) = \max\{X_i^t \in \mathcal{X} | I(c_i, q^t) = 1, t \in [s, e]\}$.

The probability of collision for a trajectory $\mathcal{T}^{s:e}$ in the time horizon from $[s, e]$ given the observations O is

$$\mathbb{E}(R(\mathcal{T}^{s:e})|O) = 1 - \prod_{c_i \in C} \prod_{t=s}^{e} \mathbb{P}(X_i^t = 0|O) I(c_i, q^t), \tag{2}$$

where the equality is due to the independence assumption on the random variables for cell occupancy.

III. Approach

In this section we propose an informative trajectory planning method for the ego vehicle to explore the environment.

A. Informative Trajectory Planning Framework

Calculating Equation (1) requires considering the different outcomes of observations at each step of the planning horizon N, i.e., O^{-nT} for all $n \in \{1, \ldots, N\}$. Therefore, the calculation of risk for a candidate trajectory is computationally expensive. A conventional way of approximating the risk at time 0 in Equation (1) along the planning horizon is by approximating $\mathbb{E}\big(\mathbb{E}(R(\mathcal{T}^{nT:(n+1)T})|O^{-nT})\big)$ with $\mathbb{E}(R(\mathcal{T}^{nT:(n+1)T})|O^{-0})$ using the observations until the current time step 0, i.e., O^{-0} [19]. Observe that to plan a safe trajectory based on O^{-0} one must assume all occluded regions are occupied [7]. This results in overestimating the risk of a trajectory in the future time steps and consequently conservative trajectories.

While planning trajectories, the ego vehicle can opt for the trajectory that gathers information on the state of cells in occluded regions. These observations improve the ego vehicle's understanding of the risks of collision in the future.

The Shannon entropy, denoted by $\mathcal{H}$, [25] is a measure of the uncertainty of a random variable given observations O. The change in Shannon entropy with observations until time s represents the information gain on the risk random variable, i.e.,

$$\Delta\mathcal{H}(R(\mathcal{T}^{nT:(n+1)T})|O^{-s}) = \mathcal{H}(R(\mathcal{T}^{nT:(n+1)T})|O^{-0}) - \mathcal{H}(R(\mathcal{T}^{nT:(n+1)T})|O^{-s}).$$

Note that observations on the cells that collide with $\mathcal{T}^{nT:(n+1)T}$ at a time closer to nT provides a better understanding of the state of the cells in time horizon $[nT, (n+1)T]$. However, early observations allow the ego vehicle to adjust its trajectory and minimize the risk of collision. Therefore, we evaluate the total change in the uncertainty of $R(\mathcal{T}^{nT:(n+1)T})$ with a linear combination of the information gain at different time steps, i.e.,

$$\Delta\mathcal{H}(R(\mathcal{T}^{nT:(n+1)T})) = \sum_{s=0}^{nT} w_s \Delta\mathcal{H}(R(\mathcal{T}^{nT:(n+1)T})|O^{-s}), \tag{3}$$

where w_s are user-defined weights representing the importance of making observations at different time steps.

Informative Trajectory Planner: Now we propose our informative trajectory planner, which minimizes the cost and the expected risk, and maximizes the information gain of the trajectory of ego vehicle. Given a set of feasible candidate trajectories in Ω, we evaluate the following objective:

$$\min_{\mathcal{T}} \Bigg(J(\mathcal{T}^{0:NT}, \mathcal{T}^{0:NT}_{\text{nom}}) + \sum_{n=0}^{N} \rho_n \mathbb{E}(R(\mathcal{T}^{nT:(n+1)T})|O^{-0}) - \sum_{n=0}^{N} \beta_n \Delta\mathcal{H}(R(\mathcal{T}^{nT:(n+1)T})) \Bigg), \tag{4}$$

where $\beta_n \geq 0$ are user-defined parameters representing the importance of information gain at different time horizons. The proposed method plans trajectories in a receding horizon manner. In other words, at the current time step $t = 0$, the ego car plans a trajectory for the next $\{0, \ldots, N\}$ planning horizon, and updates its plan at $t = T$ for the planning horizon $\{1, \ldots, N+1\}$. Maximizing the information gain in Equation (4) improves the estimate of the risk for the planning interval starting at time $t = T$, i.e., $\sum_{n=1}^{N+1} \rho_n \mathbb{E}(R(\mathcal{T}^{nT:(n+1)T})|O^{-T})$ which may result in less conservative trajectories in future planning horizons. In Equation (4), we consider a set of candidate trajectories with different lateral deviation from the nominal trajectory and different speed profiles. By deviating from the nominal trajectory, the ego vehicle is able to make observation on the future risks. A method to generate the set of candidate trajectories is the Frenet Planner [26].

Now we make the following observation on the safety of the trajectory generated by the proposed algorithm:

Lemma 1. Consider an ego vehicle starting at a collision free state. The output of the informative trajectory planner is a collision free trajectory at any time step.

The proof of the lemma comes from the definition of Ω. The trajectory given by the informative planner is in the set Ω, therefore, there is a safe trajectory starting at q^t for any time step $t \in [0, T]$. In the next planning step at time $t = T$, the ego vehicle will either take the safe trajectory to avoid collision or plans a trajectory in Ω which ensures safety for $[T, 2T]$. Then the ego vehicle starting at a collision free state will not collide while following the trajectories provided by the informative planner.

Note that calculating Equation (3) requires considering different outcomes of observations in future steps and it is computationally expensive. Therefore, in the next sections we provide a method to approximate the information gain.

B. Region to Observe

The probability that c_j in G_v is not occupied at time t given observations O^{-s} is

$$\mathbb{P}(X_j^t = 0|O^{-s}) = \Pi_{c_i \in G_v} \big[1 - \Pr{}^v_{i,j}(t-s)\mathbb{P}(X_i^s = 1|O^{-s})\big].$$

The occupation probabilities for G_p are developed in the same manner and are omitted for brevity.

Let $\ell_{i,j}^{s,t} = \Pi_{c_k \in G_v, k \neq i}[1 - \Pr{}^v_{k,j}(t-s)\mathbb{P}(X_k^s = 1|O^{-s})]$ be the probability that cell c_j is not occupied at time t given that no other cell c_k, excluding c_i, moves to c_j by time $t \geq s$.

Now we establish the following result on the information gain for the cell c_i given an observation on c_j at time s.

Lemma 2. The information gain on a cell c_i at time t given an observation on cell c_j at time s is a monotonically increasing function of transition probability $\Pr{}^v_{j,i}(t-s)$.

Proof. The information gain on cell c_i at time t is as follows:

$$\mathcal{H}(X_j^t) - \mathcal{H}(X_j^t|X_i^s) = -\sum_{z\in\{0,1\}} \mathbb{P}(X_j^t = z)\log(\mathbb{P}(X_j^t = z)) + \sum_{k\in\{0,1\}} \mathbb{P}(X_i^s = k)\Bigg[\sum_{z\in\{0,1\}} \mathbb{P}(X_j^t = z|X_i^s = k) \log(\mathbb{P}(X_j^t = z|X_i^s = k))\Bigg]. \tag{5}$$

Observe that $\mathbb{P}(X_j^t = 0|O^{-s}) = \ell_{i,j}^{s,t}(1 - \Pr{}^v_{i,j}(t-s)\mathbb{P}(X_i^s = 1|O^{-s}))$, and consequently $\mathbb{P}(X_j^t = 0|X_i^s = 1) = \ell_{i,j}^{s,t}(1 - \Pr{}^v_{i,j}(t-s))$ given that we observe an agent at X_i at time s. Taking the derivative and simplifying the equations, the derivative of the information becomes

$$-\ell_{i,j}^{s,t}\mathbb{P}(X_i^s = 1)\log\frac{1 - \ell_{i,j}^{s,t} - \ell_{i,j}^{s,t}\Pr{}^v_{i,j}(t-s)\mathbb{P}(X_i^s = 1)}{\ell_{i,j}^{s,t} - \ell_{i,j}^{s,t}\Pr{}^v_{i,j}(t-s)\mathbb{P}(X_i^s = 1)} + \ell_{i,j}^{s,t}\mathbb{P}(X_i^s = 1)\log\frac{1 - \ell_{i,j}^{s,t} - \ell_{i,j}^{s,t}\Pr{}^v_{i,j}(t-s)}{\ell_{i,j}^{s,t} - \ell_{i,j}^{s,t}\Pr{}^v_{i,j}(t-s)}.$$

To prove the result, it suffices to show that the derivative is positive for any $\Pr{}^v_{i,j}(t-s) \in [0, 1]$. After some simplifications to the inequality above, we arrive at following condition

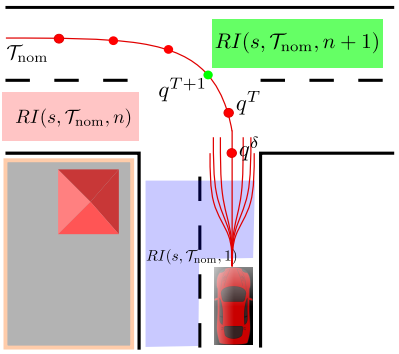

Fig. 3: Three regions of interest at time s for the future trajectory of the ego vehicle: the immediate surroundings (blue), oncoming traffic in the lane to be crossed (red), and traffic in the destination lane (green).

$\Pr^v_{i,j}(t-s)(1-\mathbb{P}(X_i^s=1)) + 2\ell^{s,t}_{i,j}\Pr^v_{i,j}(t-s)\mathbb{P}(X_i^s=1)(1-\Pr^v_{i,j}(t-s)) \geq 0$, which is correct for any values of $\ell^{s,t}_{i,j}, \mathbb{P}(X_i^s=1)$, and $\Pr^v_{i,j}(t-s) \in [0,1]$. □

Lemma 2 shows observing a cell c_i^s at time s with higher transition probability $\Pr^v_{i,j}(t-s)$ increases the information gain on the state of c_j at time t. The monotonicity of information gain with respect to the transition probabilities can be further leveraged to limit the regions to observe.

Definition 2 (Region of Interest)**.** Given a trajectory $\mathcal{T}$ and a planning step $n \in \{0,\dots,N\}$, the region of interest $RI(s,\mathcal{T},n)$ at time $s < (n+1)T$ is the set of cells from which an actor can reach a collision state with the vehicle on trajectory $\mathcal{T}$ in time interval $[nT,(n+1)T]$. Formally, this is the set $RI(s,\mathcal{T},n) = \{c_i \in C | \exists c_j \in RS(c_i, t-s), I(c_j,q^t)=1, t\in[nT,(n+1)T]\}$.

C. Information Gain on Individual Cells

Next we connect the information gain on collision risk in future steps to individual cells.

Lemma 3. The information gain on a cell c_j at time t given an observation on c_i at time s is a concave function of $\mathbb{P}(X_i^s=1)$.

Proof. Observe that the latter four terms in the right hand side of Equation (5) are linear in $\mathbb{P}(X_i^s=1)$. Therefore, it suffices to show that the first two terms in the right hand side are concave functions in $\mathbb{P}(X_i^s=1)$. Note that

$$\partial^2\mathcal{H}(X_j^t)/\partial\mathbb{P}(X_i^s=1)^2 = -\frac{\ell^{s,t}_{i,j}\Pr^v_{i,j}(t-s)^2}{1-\Pr^v_{i,j}(t-s)\mathbb{P}(X_i^s=1)} - \frac{\left(\ell^{s,t}_{i,j}\right)^2\Pr^v_{i,j}(t-s)^2}{1-\ell^{s,t}_{i,j}+\ell^{s,t}_{i,j}\Pr^v_{i,j}(t-s)\mathbb{P}(X_i^s=1)}.$$

Note that $\partial^2\mathcal{H}(X_i^t)/\partial\mathbb{P}(X_i^s=1)^2 \leq 0$ for all values of $\mathbb{P}(X_j^s=1)$, therefore, the result follows immediately. □

Lemma 3 shows that the information gain is a concave function in $\mathbb{P}(X_i^s=1)$. Finding $\mathbb{P}(X_i^s=1)$ with the maximum information gain requires finding the solution to equation $\partial\mathcal{H}(X_j^t)/\partial\mathbb{P}(X_i^s=1)=0$ and we omit it in this paper. The range of the $\mathbb{P}(X_i^s=1)$ that maximizes the information gain for different values of $\ell^{s,t}_{i,j}$ and $\Pr^v_{i,j}(t-s)$ is $[0.41, 0.61]$. Hence, observing the cells at time s with higher entropy carries more information on the state of the cells colliding with the trajectory. By Lemma 3, we can approximate $\mathcal{H}(R(\mathcal{T}^{nT:(n+1)T}|O^{-s}))$ with $\sum_{c_j\in RI(s,\mathcal{T},n)}\mathcal{H}(X_j^s|O^{-s})$ which is the total information gain on individual cells in the region of interest at time s. By substituting in Equation (3), we approximate the information gain as follows:

$$\Delta\mathcal{H}(R(\mathcal{T}^{nT:(n+1)T})) \approx \sum_{s=0}^{nT} w_s\Big[\mathcal{H}(R(\mathcal{T}^{nT:(n+1)T})|O^{-0}) - \sum_{c_j\in RI(s,\mathcal{T},n)}\mathcal{H}(X_j^s|O^{-s})\Big].$$

Observe that by Lemma 2 we find the set of locations to observe and by the result of Lemma 3 we approximate the information gain on risk of collision over the the cells in the region of interest. To reduce the computational cost of calculating the information gain for time horizon $[0,NT]$, we evaluate each candidate trajectory's information gain based the cells in the region of interest of the nominal trajectory (see Figure 3). Given predictions of occupancy for the two grids at time s, we unify them to a single grid by mapping the locations in G_v to G_p. For the nominal trajectory $\mathcal{T}_{\text{nom}}$, we find the region of interest at each time step s, i.e., $RI(s,\mathcal{T}_{\text{nom}},n)$. The region of interest and the unified occupancy grid are inputs to the module that calculates the probability of visibility for the cells in $RI(s,\mathcal{T}_{\text{nom}},n)$ for different lateral deviations from the nominal trajectory. The visibility module outputs a set of grids each representing the visibility of trajectories of certain lateral deviations.

IV. Simulation Results

In this section, we evaluate the performance of the proposed informative trajectory planner in different scenarios. The simulations are performed in the CARLA [27] environment using a AMD Ryzen7 2700 CPU and 1080 Nvidia GPU. Feasible trajectories are generated using a Frenet Planner with a maximum velocity of 8 m/s, a lateral step size of 0.5m, and different speed profiles over 2-4 seconds. We use $\rho_n = 10^3$ and $\beta_n = 10^2$ for all $n \in \{1,\dots,N\}$. Finally the information gain weight w_i is the probability of the ego vehicle making the observation, $\prod_k \mathbb{P}(X_k^t=0)$ for cells c_k in the ray cast from the ego vehicle to cell c_i. The cost function J for these experiments is a linear combination of the lateral deviation from the nominal trajectory and its length.

a) Evaluation of Informative Planning: In the first experiment, we evaluate the performance of the proposed informative trajectory planner in a parking lot scenario illustrated in Figure 4. The visibility of the ego vehicle is obstructed by the vehicles highlighted with purple rectangles. We compare the results to the probabilistic occlusion-aware trajectory planner proposed in [19]. The ego vehicle and the region of interest are highlighted by green and red rectangles, respectively. Moving vehicles are initiated at the

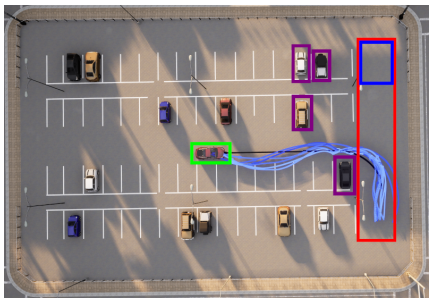

Fig. 4: A comparison of a centre-line trajectory (black) with representative variations of the information gain trajectory (shades of blue). Detecting an agent in the end file (red zone) can cause the Information Gain trajectory to move towards the nominal resulting in the different trajectories shown.

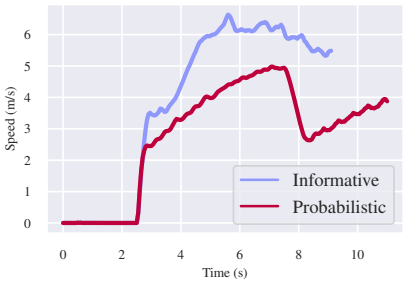

Fig. 5: Comparison of the proposed informative algorithm vs. a probabilistic planner [19].

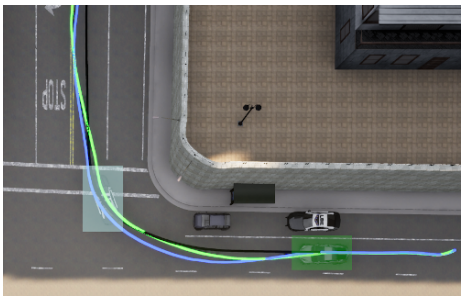

Fig. 6: Comparison of information gain vs. maximizing visibility. When the ego vehicle is following another vehicle that blocks future observations, the Information Gain (green line) method returns to nominal (black line), whereas the visibility maximization (blue line) continues to swerve.

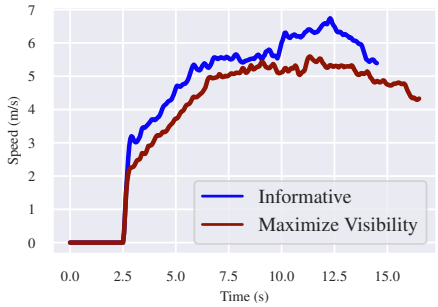

Fig. 7: Comparison of the proposed informative algorithm vs. maximizing the visibility.

region outlined with blue (see Figure 4) according to a Poisson process with arrival rate $0.3/s$. Trajectories from the proposed informative planner are shown in shades of blue, the probabilistic method in black. To maximize the information gain, the vehicle swerves to right initially, but as it approaches the corner, it switches focus and swerves to the left to better observe the lower right corner. On average, after 40 experiments, the proposed method observed the actors originating from the blue region 0.4 seconds earlier than the probabilistic method. The same experiment was performed with no additional actors resulting in the speed profile shown in Figure 5. The probabilistic method slows to ensure safety as it has limited visibility beyond the turn. However, the informative approach, by adjusting its trajectory, is able to observe locations earlier and proceed at a consistent speed.

b) Information Gain vs. Visibility: In this experiment, compare our method to that proposed in [21] where the planner works to minimize occlusions. Figure 6 shows the scenario where the ego vehicle (highlighted with a green rectangle) is making a right turn while its vision is obstructed by the static vehicles and the walls. There is an actor (the white rectangle) ahead of the ego vehicle performing the same maneuver. The visibility maximization method, shown as a blue trajectory, minimizes the total occlusion without regard to the movement of the other actors. On the other hand, the trajectory of the informative planner, shown in green, initially deviates slightly from the nominal trajectory since there is a probability that the white vehicle does not turn right. However, after observing that the white vehicle will block any future observations of the region of interest, the information gain method converges to the nominal trajectory as there is no advantage in deviating. If the white vehicle is not present then the information gain method captures the behaviour of algorithms that minimize occlusions. Figure 7 shows the speed profile of the proposed algorithm and the algorithm to maximize the visibility. Observe that the two methods have similar speed profiles; however, due to larger deviation from the nominal trajectory and longer paths, the method to minimize occlusions is slower to reach the goal.

V. Conclusion and Future Work

In this paper, we introduced the problem of informative trajectory planning for ego vehicles to observe the potential collision risk originating from occluded areas. Then we provided an approximation of the information gain and proposed an informative trajectory planning framework. Finally, we illustrated the performance of the proposed algorithm in a set of experiments, comparing our results with an algorithm that is limited to velocity changes along the nominal path, and one that alters the vehicle's trajectory to minimize occluded regions. For future work, we would like to investigate the performance of the informative motion planning framework in scenarios where multiple vehicles perform informative maneuvers.

REFERENCES

[1] P. F. Orzechowski, A. Meyer, and M. Lauer, "Tackling occlusions & limited sensor range with set-based safety verification," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018, pp. 1729–1736.

[2] D. Fehr, W. J. Beksi, D. Zermas, and N. Papanikolopoulos, "Occlusion alleviation through motion using a mobile robot," in *2014 International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 3179–3184.

[3] E. Gonzalez Debada, A. Ung, and D. Gillet, "Occlusion-Aware Motion Planning at Roundabouts," *Transactions on Intelligent Vehicles*, 2020.

[4] 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.

[5] S. Hoermann, F. Kunz, D. Nuss, S. Renter, and K. Dietmayer, "Entering Crossroads with Blind Corners. A Safe Strategy for Autonomous Vehicles," in *2017 Intelligent Vehicles Symposium (IV)*. IEEE, 2017, pp. 727–732.

[6] R. De Iaco, S. L. Smith, and K. Czarnecki, "Universally safe swerve maneuvers for autonomous driving," *IEEE Open Journal of Intelligent Transportation Systems*, 2021.

[7] Y. Nager, A. Censi, and E. Frazzoli, "What lies in the shadows? safe and computation-aware motion planning for autonomous vehicles using intent-aware dynamic shadow regions," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 5800–5806.

[8] A. Lawitzky, A. Nicklas, D. Wollherr, and M. Buss, "Determining states of inevitable collision using reachability analysis," in *2014 International Conference on Intelligent Robots and Systems (IROS)*. IEEE/RSJ, 2014, pp. 4142–4147.

[9] F. Damerow, T. Puphal, Y. Li, and J. Eggert, "Risk-based driver assistance for approaching intersections of limited visibility," in *2017 International Conference on Vehicular Electronics and Safety (ICVES)*. IEEE, 2017, pp. 178–184.

[10] M. Lee, K. Jo, and M. Sunwoo, "Collision risk assessment for possible collision vehicle in occluded area based on precise map," in *2017 20th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2017, pp. 1–6.

[11] L. Wang, C. F. Lopez, and C. Stiller, "Generating Efficient Behaviour with Predictive Visibility Risk for Scenarios with Occlusions," in *2020 23rd International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2020, pp. 1–7.

[12] M. Koschi and M. Althoff, "Set-based Prediction of Traffic Participants Considering Occlusions and Traffic Rules," *Transactions on Intelligent Vehicles*, 2020.

[13] G. Neel and S. Saripalli, "Improving bounds on occluded vehicle states for use in safe motion planning," in *2020 International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE, 2020, pp. 268–275.

[14] M.-Y. Yu, R. Vasudevan, and M. Johnson-Roberson, "Risk assessment and planning with bidirectional reachability for autonomous driving," in *2020 International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 5363–5369.

[15] S. Hoermann, M. Bach, and K. Dietmayer, "Dynamic occupancy grid prediction for urban autonomous driving: A deep learning approach with fully automatic labeling," in *2018 International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 2056–2063.

[16] C. Fulgenzi, A. Spalanzani, and C. Laugier, "Dynamic obstacle avoidance in uncertain environment combining pvos and occupancy grid," in *2007 International Conference on Robotics and Automation (ICRA)*. IEEE, 2007, pp. 1610–1616.

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

[18] X. Lin, J. Zhang, J. Shang, Y. Wang, H. Yu, and X. Zhang, "Decision Making Through Occluded Intersections for Autonomous Driving," in *2019 Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2019, pp. 2449–2455.

[19] S. G. McGill, G. Rosman, T. Ort, A. Pierson, I. Gilitschenski, B. Araki, L. Fletcher, S. Karaman, D. Rus, and J. J. Leonard, "Probabilistic risk metrics for navigating occluded intersections," *Robotics and Automation Letters*, vol. 4, no. 4, pp. 4322–4329, 2019.

[20] T. Nyberg, C. Pek, L. Dal Col, C. Norén, and J. Tumova, "Risk-aware Motion Planning for Autonomous Vehicles with Safety Specifications," in *2021 Intelligent Vehicles Symposium (IV)*. IEEE, 2021.

[21] J. Higgins and N. Bezzo, "Negotiating visibility for safe autonomous navigation in occluding and uncertain environments," *Robotics and Automation Letters*, vol. 6, no. 3, pp. 4409–4416, 2021.

[22] H. Andersen, W. Schwarting, F. Naser, Y. H. Eng, M. H. Ang, D. Rus, and J. Alonso-Mora, "Trajectory optimization for autonomous overtaking with visibility maximization," in *2017 20th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2017, pp. 1–8.

[23] N. Buckman, A. Pierson, S. Karaman, and D. Rus, "Generating Visibility-Aware Trajectories for Cooperative and Proactive Motion Planning," in *2020 International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 3220–3226.

[24] G. Jahangirova, A. Stocco, and P. Tonella, "Quality Metrics and Oracles for Autonomous Vehicles Testing," in *2021 14th Conference on Software Testing, Verification and Validation*. IEEE, 2021, pp. 194–204.

[25] C. E. Shannon, "A Mathematical Theory of Communication," *The Bell system technical journal*, vol. 27, no. 3, pp. 379–423, 1948.

[26] A. Sakai, D. Ingram, J. Dinius, K. Chawla, A. Raffin, and A. Paques, "PythonRobotics: a Python code collection of robotics algorithms," *arXiv preprint arXiv:1808.10703*, 2018.

[27] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An Open Urban Driving Simulator," in *Proceedings of the 1st Annual Conference on Robot Learning*, 2017, pp. 1–16.